The Accelerator Neutrino Neutron Interaction Experiment: Design, Construction and Preparation for First Physics

Marcus O'Flaherty
Dept of Particle Physics and Particle Astrophysics
University of Sheffield

A thesis submitted for the degree of
PhD in Physics

May 2020

Foreword

This thesis covers work performed for the Accelerator Neutrino Neutron Interaction Experiment (ANNIE) between 2015 and 2020. ANNIE is a beamline neutrino experiment located at the Fermi National Accelerator Laboratory in Chicago. Its primary physics goals are the measurement of neutrino interaction cross-sections on oxygen, and a measurement of the multiplicity of final state neutrons from these interactions, both as functions of final state kinematics. In order to achieve these goals ANNIE will be pioneering the use of Large Area Picosecond Photodetectors and the use of gadolinium doping to enhance neutron visibility in a water Cherenkov detector.

Over the time period of this thesis the ANNIE detector was constructed, an initial 'Phase I' background run was completed, the detector was decommissioned and then substantially upgraded in preparation for the main 'Phase II' physics run. At the time of writing ANNIE is being re-commissioned, with Phase II data taking anticipated to begin in 2021.

During this time I spent 2½ years on-site contributing to almost all aspects of hardware work, while performing software development for the simulation and analysis frameworks.

As my first software task I developed Monte Carlo simulations of the detector response, based on the WCSim framework. While this framework provided a starting point for the tank simulation significant modifications were required to implement the complete ANNIE detector system. This included a rewrite of the geometry, along with modifications to the photosensor definitions, hit collections, digitization, triggering, input and output interfaces, etc. To take advantage of improved modelling of neutrons in Geant4.10 the physics lists were also reimplemented in the new version format. As part of this work I determined performance metrics based on the output files, calculating kinematic acceptance, detector efficiency and expected event rates.

My contributions on reconstruction focused primarily on the MRD, for which I implemented two track reconstruction algorithms, specialised for high and low energy muons. The high energy algorithm achieves >90% efficiency for tracks with at least 8 hits, and is robust against dead paddles and unrelated hits. The low energy algorithm performs a search in combination with tank information to provide sensitivity down to a single MRD hit.

Early data handling work for the Phase I background run included writing the code for handling TDC hits and performing studies on neutron detection without the requirement of multi-hit coincidence. As the collaboration moved to a new framework for the Phase II analysis I performed a substantial amount of work to develop the codebase, implementing tools to perform tasks such as file IO, data parsing, waveform processing, event building, timestamp alignment, time clustering, hit residual calculation, event visualisation etc. I helped develop the python API for integrating multi-language tools, and ported the framework into a Singularity container to provide a self-contained environment that could run on scientific computing clusters such as FermiGrid or the Open Science Grid. Application specific studies included toolchains to optimise MRD discriminator thresholds and to extract PMT gains from LED calibration data.

On more hardware oriented aspects I have been in charge of high voltage supplies throughout my time on ANNIE. This ranged from system design to testing, assembly, installation and maintenance. The Phase I system provided 200 output channels and ran with 98% uptime over the 2 year data taking period. I then upgraded the system to 416 channels in preparation for Phase II, where it has been running with >99% uptime since data taking started. This work included the development of custom software for hardware control, data logging, local and web-based monitoring.
As part of the Phase II upgrade I performed work on the MRD refurbishment, repairing and installing the replacement scintillator paddles and assembling the electronics racks to acquire TDC data. I also carried out the initial ex-situ calibration of water tank PMTs, developing a standalone test station and associated data acquisition software. I performed testing and analysis of the data to extract voltage vs gain curves, using this information to determine initial setpoints for PMT voltages.

I would like to express my gratitude to my supervisor, Dr Matthew Malek, to my collaborators on ANNIE, and to the staff at the Fermilab facility for all the support and encouragement they gave to me over the course of my studies.

Contents

1 Introduction and Background **1**

1.1 Neutrinos in The Standard Model . . . 1

1.2 Majorana or Dirac . . . 6

1.3 Fundamental Neutrino Interactions . . . 10

1.4 Neutrino-Nucleus Interactions . . . 25

1.5 The Importance of Neutrino Cross-Sections . . . 36

1.6 Measurement of Neutrino Cross-Sections . . . 38

1.7 The Importance of Neutron Multiplicity . . . 42

1.8 Status of Neutrino Cross Section Measurements . . . 44

1.9 Status of Neutrino Multiplicity Measurements . . . 47

2 The ANNIE Experiment **48**

2.1 Forward Muon Veto (FMV) . . . 52

2.2 Water Cherenkov Detector . . . 54

2.3 Muon Range Detector . . . 62

2.4 Outer Detector . . . 64

2.5 Data Acquisition System (DAQ) . . . 66

2.6 Control and Readout . . . 68

2.7 High Voltage Supplies . . . 69

2.8 Gadolinium Doping . . . 70

2.9 Chemical Compatibility . . . 74

2.10 Water Quality . . . 78

2.11 Large Area Picosecond Photodetectors (LAPPDs) . . . 82

3 Calibration **89**

3.1 PMTs . . . 89

3.2 LAPPDs . . . 99

3.3 MRD . . . 100

3.4 FMV 103

4 The Fermilab Booster Neutrino Beam **104**

4.1 Proton Beam Generation 104

4.2 Neutrino Beam Generation 105

4.3 Beam Monitoring 108

5 Monte Carlo Simulations **109**

5.1 Hadron Production 109

5.2 Neutrino Flux at the ANNIE Detector 114

5.3 Neutrino Interactions 116

5.4 Background Sources 118

5.5 Detector Simulation 123

6 Reconstruction **128**

6.1 Time Clustering 128

6.2 Hit Cleaning 129

6.3 Vertex Reconstruction 131

6.4 MRD Track Reconstruction 137

6.5 Tank Energy Reconstruction 141

6.6 Event Classification 143

7 Analysis **146**

7.1 Analysis Strategy 146

7.2 Uncertainties 148

8 The Future **153**

8.1 Next Steps 153

8.2 Additional Opportunities 154

9 Conclusion **156**

References **157**

1

Introduction and Background

1.1 Neutrinos in The Standard Model

The Standard Model of particle physics outlines the set of quantum fields that describe the physical world at the subatomic scale. The evolution of these fields is most commonly calculated via Lagrangian mechanics, which describes the equations of motion as a set of Euler-Lagrange equations. Denoting the Lagrangian density $\mathcal{L}$, the Euler-Lagrange equation for a field Ψ is:

$$\partial_\mu \frac{\delta L}{\delta(\partial_\mu \Psi)} = \frac{\delta \mathcal{L}}{\delta \Psi} \tag{1.1}$$

where the use of δ indicates derivative with respect to the field function.
Perhaps the most important principle of the Standard Model is one of symmetry. As described by Noether's theorem[1], for every symmetry of the Lagrangian there is a conserved current and charge, and within the Standard Model these are the particle currents and their charges. Finding the fields to describe the physical particles then comes down to finding a Lagrangian that satisfies the corresponding symmetries.

Since its development in the 1970s the Standard Model has seen immense success in describing and predicting the properties and interactions of particles, and has upheld incredibly scrutinous testing. Searches for new particles and processes in collider experiments such as ATLAS[2] and CMS[3] have turned up nothing unexpected, nor have searches for rare processes in experiments such as LUX[4], Xenon1T[5], or Super-Kamiokande[6]. Precision measurements such as those of the electron dipole moment[7], lepton universality[8], and muon decay distributions[9], among many others[10], have all returned results consistent with Standard Model predictions. But while it has proven incredibly precise in predictions within its scope, the Standard Model is known to be incomplete. A few key problems highlight its shortcomings:

1. Gravity

   The Standard Model describes three of the four fundamental forces - the strong, weak, and electromagnetic forces - but gravity, whose influence is undetectable on the small scales of particle interactions, is left out. While the theory of General Relativity has proven just as successful in describing gravitational interactions the two frameworks describe nature in fundamentally different ways, and their unification is an as-yet unsolved mystery.

2. Dark Energy

   An example of the jarring boundary between these frameworks comes from the accelerating expansion of the universe. The Standard Model of cosmology is the ΛCDM model; a model requiring both cold dark matter and a positive cosmological constant[11]. This cosmological constant represents the energy density of the vacuum, a source of negative pressure that overcomes the attractive force of gravity and drives the universe into accelerating expansion. Quantum Field Theory tells us such a vacuum energy indeed exists; the quantum vacuum is a sea of particle-antiparticle pairs, continually popping into and out of existence, and therefore has a ‘zero point energy’. But attempts to calculate the corresponding cosmological constant result in overshooting the required value by a factor of 10^{55}[12]!

3. Dark Matter

   Cosmological models not only require dark energy, but also dark matter; a non-relativistic, non-interacting form of matter necessary for describing kinematics at and above galactic scales. Observations indicate this unknown source of gravitational attraction could be a gas of weakly interacting massive particles (WIMPs), but initial efforts at fitting neutrinos into this role proved a poor match[13], leaving no suitable candidate within the Standard Model. Remaining WIMP candidates, along with other possible explanations such as axions, dark photons or even a complete dark sector[14], must reach beyond the Standard Model.

4. The Matter-Antimatter Asymmetry

   The Standard Model places matter and antimatter on very close footing; each is created, interacts, decays and annihilates almost identically, with only very minor differences. Yet these differences, where the combined symmetries of charge and parity (CP) are together violated, must account for the dominance of matter in our universe. Known sources of CP violation in the Standard Model, measured in meson mixing and decays, fall far short of the level required to generate the observed asymmetry[15].

5. Neutrinos

   Within the framework of the Standard Model neutrinos are massless, weakly interacting, left-handed fermions (or right-handed antifermions). But the discovery that neutrinos oscillate between flavour states[16] demonstrated that these are not eigenstates of the free Hamiltonian, and therefore that neutrinos are not massless. This in turn implies the existence of right-handed neutrinos. Both of these are new features that need to be included into the Standard Model, and are tantalizing indications of new physical processes at higher energy scales.

Most of these shortfalls are based on cosmological observations, and their investigation is generally theoretical, cosmological, or centre around searches for the unknown. Neutrinos are unique in offering a tangible handle on physics beyond the Standard Model. Here there are several clear avenues of investigation.
While neutrino oscillation was discovered over 20 years ago[16] the theoretical description requires knowledge of three mixing angles and two squared mass differences, and there is still work to be done in fully determining their values. Of the three mixing angles, θ_{12}, θ_{13} and θ_{23}, the former two have now been precisely measured. The mixing angle θ_{23} has been found to be very close to, but not quite 45° (global fits suggest a deviation of $\sim 3°$[14]), but degeneracies in the oscillation probability make it unclear whether $\theta_{23} > 45°$ or $< 45°$.

Neutrino oscillations are also only sensitive to neutrino mass differences; the absolute scale of the masses must be measured by other means. Although non-zero, these masses are puzzlingly small - less than 10^{-6} that of the electron. Being so small, physical processes that are influenced by neutrino masses are very well described by a massless approximation, making the measurement of neutrino masses incredibly difficult. At present only an upper limit is known[17]; a successful determination would provide great insight into the scales of new physics.
More tangible is the determination of the neutrino mass ordering. Whether neutrino masses are 'quasi-degenerate', with mass splittings much less than the lightest neutrino mass, $m_1 \simeq m_2 \simeq m_3$, or hierarchical, either with $m_1 \ll m_2 < m_3$ or with $m_3 \ll m_1 < m_2$, is also as-yet undetermined, but this could be measured by near-future long baseline experiments[18].

Finally, the very nature of neutrino masses is itself a mystery. The Standard Model is a chiral field theory; fermions are described by chirally left- or right-handed Weyl fields. Particles, representing excitations of these fields, gain mass in the coupling of left- and right-handed fields; such interactions transform excitations of a left-chiral field into those of a right-chiral one.

Denoting a chirally right-handed field Ψ and a chirally left-handed one ψ, the corresponding mass term would be $(\overline{\Psi}m\psi + \overline{\psi}m\Psi)$, where $\overline{\Psi} = \Psi^{\dagger}\gamma^{0}$ is the adjoint of Ψ.

The weak interaction, however, breaks chiral symmetry as it couples to left-handed fermions (or right-handed anti-fermions) but *not* right-handed fermions (or left-handed anti-fermions). A coupling from left- to right-chiral fields would then break local gauge invariance. To accommodate this the Standard Model introduces the Higgs mechanism; a spontaneous breaking of the electroweak symmetry that introduces a scalar field with non-zero vacuum expectation value. This field The coupling from left- to right-chiral fields may then be mediated by an excitation in the Higgs vev, which absorbs the difference in charges between the left- and right-chiral fields (the combination of fermion and Higgs excitation is a completely neutral singlet). The resulting 'Dirac mass' m_D is related to the Higgs vacuum expectation value v and the Yukawa coupling of the particle to the Higgs field Y, via $m_D = Yv/\sqrt{2}$.

A Dirac neutrino would then be a combination of two Weyl spinors $\frac{1}{2}(\Psi + \psi)$. But at the time the Standard Model was formed neutrinos were believed to be massless, so there was no mixing between Ψ and ψ. In fact, the right-chiral field ψ has no weak isospin, no hypercharge, and no strong colour, so with neither interactions nor mass it was omitted from the model. With the discovery of non-zero masses a right-handed field is now necessary, so the addition of such a field seems like the obvious choice. This does, however, raise a question: if the neutrinos gain mass through the same mechanism as the other fermions, why is their mass so much smaller? The current upper limits require Yukawa couplings to the Higgs of less than 10^{-13}!

In 1937 Ettore Majorana pre-empted this puzzle with another possibility. He demonstrated that if Ψ is chirally left-handed then the combination $\Psi^C = C\overline{\Psi}^T$ is chirally right-handed. Here C is the charge conjugation operator, which may be represented $C = i\gamma_2\gamma_0$, and the field Ψ^* is the usual complex conjugate of Ψ. With this combination a mass term may be built from just Ψ by the combination $\frac{M}{2}\left(\overline{\Psi^C}\Psi + \overline{\Psi}\Psi^C\right)$. Such 'Majorana masses' couple a field to its charge-conjugate, Ψ^C, making them forbidden for all other Standard Model particles by conservation of electric or strong charge - the right-handed neutrino represents the only allowed Majorana field. They would couple to the active states through oscillations,

It is also possible that both of these mass-generation mechanisms are present. Such right-handed Majorana neutrinos could then couple to the active left-handed states through the Higgs mechanism introducing off-diagonal entries in the mass matrix.

Once diagonalized these contributions give a compelling explanation for the discrepancy in the observed neutrino masses. Such Majorana neutrinos would also be self-conjugates - they are their own antiparticle and the mass couplings would violate lepton number conservation by 2. Both of these properties make them a natural candidate for generating the matter-antimatter asymmetry of the universe via leptogenesis[19].

The investigation of neutrinos, their oscillations, masses, and their Dirac or Majorana nature, presents a rich field of research at the forefront of tangible searches for new physics. In order to study these properties we need to understand how neutrinos interact, and how we may observe these interactions in experimental detectors. The study of these interactions is the purpose of the ANNIE experiment, and the subject of this thesis.

1.2 Majorana or Dirac

In order to give mass m to a fermion Ψ one should introduce a term to the Lagrangian that couples a chirally left-handed field Ψ_L to a chirally right-handed field Ψ_R. This may be seen by considering the chiral projection operator, $\frac{1\pm\gamma^5}{2}$, operating on the solutions of a fermionic Lagrangian containing kinetic and mass terms:

$$\mathcal{L} = \overline{\Psi}(i\partial\!\!\!/ - m)\Psi \tag{1.2}$$

$\overline{\Psi} = \Psi^\dagger\gamma^0$ represents the adjoint of Ψ. The eigenvectors of the chiral projection operator are:

$$\begin{aligned} \Psi_R &= \frac{1+\gamma^5}{2}\Psi \\ \Psi_L &= \frac{1-\gamma^5}{2}\Psi \end{aligned} \tag{1.3}$$

which have eigenvalues -1 (Ψ_L) and +1 (Ψ_R) respectively. Writing the general wavefunction as a sum of these eigenvectors, Equation 1.2 becomes:

$$\begin{aligned} \mathcal{L} &= (\overline{\Psi_L} + \overline{\Psi_R})(i\partial\!\!\!/ - m)(\Psi_L + \Psi_R) \\ &= \overline{\Psi_L}(i\partial\!\!\!/)\Psi_L + \overline{\Psi_R}(i\partial\!\!\!/)\Psi_R \\ &+ \overline{\Psi_L}(i\partial\!\!\!/)\Psi_R + \overline{\Psi_R}(i\partial\!\!\!/)\Psi_L \\ &- \overline{\Psi_L}m\Psi_L - \overline{\Psi_R}m\Psi_R \\ &- \overline{\Psi_L}m\overline{\Psi_R} - \overline{\Psi_R}m\Psi_L \end{aligned} \tag{1.4}$$

We can use the commutation properties of the γ matrices to show that $\overline{\Psi_L} = \overline{\Psi}P_R$, and thus $\overline{\Psi_L}\Psi_L = \overline{\Psi}P_RP_L\Psi = 0$, eliminating same-chirality mass terms, while $\overline{\Psi_L}\gamma^\mu\Psi_R = \overline{\Psi}P_R\gamma^\mu P_R\Psi = \overline{\Psi}P_RP_L\gamma^\mu\Psi = 0$, eliminating cross-chiral kinetic terms. This reduces Equation 1.4 to:

$$\mathcal{L} = i\overline{\Psi_L}\partial\!\!\!/\Psi_L + i\overline{\Psi_R}\partial\!\!\!/\Psi_R - m(\overline{\Psi_L}\Psi_R + \overline{\Psi_R}\Psi_L) \tag{1.5}$$

This demonstrates that a mass term couples left-chiral and right-chiral fields, and that for neutrinos to be massive both a chirally left-handed and right-handed field are necessary. It is worth noting that, other than via this coupling, Ψ_L and Ψ_R may otherwise be completely distinct fields.

1.2.1 Dirac Mass

If we try to identify the field Ψ as a fermion field in the weak interaction basis, one sees that such terms would not conserve weak isospin; left-handed fermions are described by electroweak doublets (weak isospin 1/2), while right-handed fermions are electroweak singlets (weak isospin

0), thus the left- and right-chiral fermion fields transform differently under gauge transformations. Denoting the weak generators T and hypercharge generator Y and gauge transformations $\alpha(x)$ and $\beta(x)$, we have:

$$\begin{aligned} \Psi'_L &\rightarrow \Psi_L = e^{i\alpha(x)T+i\beta(x)Y}\Psi_L \\ \Psi'_R &\rightarrow \Psi_R = e^{i\beta(x)Y}\Psi_R \end{aligned} \tag{1.6}$$

One can see that the product $m\overline{\Psi}\Psi$ is not gauge invariant. In order to allow masses for the fermions in the Standard Model, we must introduce Spontaneous Symmetry Breaking through the non-zero expectation value of the Higgs field. Above the electroweak phase transition the $SU(2)_L \times U(1)_Y$ symmetry is preserved, but below it weak isospin conservation may be violated by interactions with the charged vacuum. This is achieved via the introduction of an isospin doublet comprising a scalar field $H(x)$ with vacuum expectation value v:

$$\phi = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ v + H(x) \end{pmatrix} \tag{1.7}$$

$$\widetilde{\phi}^c = -i\tau_2\phi^* = \frac{1}{\sqrt{2}}\begin{pmatrix} v + H(x) \\ 0 \end{pmatrix} \tag{1.8}$$

This introduces several new terms to the Lagrangian including couplings between the fermions and this field:

$$-\sum_{\alpha\beta}[\ \overline{\Psi}^L_\alpha Y_{\alpha\beta}\phi\Psi^R_\beta + \text{h. c.}\ ] - \sum_{\alpha\beta}[\ \overline{\Psi}^L_\alpha Y_{\alpha\beta}\widetilde{\phi}^c\Psi^R_\beta + \text{h. c.}\ ] \tag{1.9}$$

Here 'h.c.' denotes the Hermitian conjugate and $Y_{\alpha\beta}$ are the matrices of Yukawa couplings. We now identify that each fermion has a mass related to its coupling to the Higgs field; $m = Y_{\alpha\beta}v$.

At the time the Standard Model was formed neutrinos were believed to be massless. Given the parity asymmetry of weak interactions, along with zero electric charge and zero mass, all observable neutrinos would necessarily be part of left-handed weak doublets. A right-handed neutrino field would be completely devoid of interactions, and there would be no mechanism by which they could be generated; having no meaningful presence they were omitted from the Standard Model. Without a right-handed counterpart, terms of the form in Equation 1.9 are of course absent for neutrinos.

The discovery of non-zero neutrino masses can easily be accommodated into the Standard Model by introducing the missing right handed neutrinos, allowing interactions with the Higgs field as in Equation 1.9. Mass generated this way is called the 'Dirac mass'. Note that $Y_{\alpha\beta}$ need not be diagonal, in which case weak interactions would produce a superposition of mass eigenstates, resulting in flavour mixing. The Higgs mechanism provides a straightforward way to give masses to neutrinos in the Standard Model, but the smallness of the observed masses would require incredibly small Yukawa couplings $Y < 10^{-13}$; although not impossible this would suggest physics of some higher energy scale is involved.

1.2.2 Majorana Mass

The previous subsection began by saying that a mass term requires the coupling of left- and right-handed chiral fields. These fields are generally distinct; for all Standard Model particles the Dirac masses couple two different fields. It is possible, however, to build a mass term from a single field by realising that if Ψ_L represents a chirally left-handed field then $\Psi_L^c = -i\gamma^2\Psi_L^*$ is chirally right-handed. The resulting term would be:

$$\mathcal{L}_M = -\frac{m_L}{2}\overline{\Psi_L^c}\Psi_L + \overline{\Psi_L}\Psi_L^c \tag{1.10}$$

Such Majorana masses couple a field to its charge-conjugate, thereby violating conservation of all charges carried. To give Majorana masses to the Standard Model neutrinos would require either the introduction of a new Higgs-like triplet field with hypercharge +2 and weak isospin -1 (Equation 1.11), or an effective operator involving two Higgs doublets arranged to transform as a triplet (Equation 1.12).

$$H_{I3\neq 0}\overline{\Psi_L^c}\Psi_L \tag{1.11}$$

$$H_{SM}H_{SM}\overline{\Psi_L^c}\Psi_L \tag{1.12}$$

For a right-handed neutrino field, which carries no charge quantum numbers, the transition from field to charge conjugate results in no charge violation. In all Standard Model interactions lepton number itself is conserved, but this may simply be an accidental symmetry of the SM interactions. If so, terms of the form in Equation 1.10 would be permitted, and observation of lepton number violating processes would be direct evidence of Majorana masses for right-handed neutrinos.

In general these two mass generation mechanisms may both be present in the Lagrangian. Active left-handed neutrinos would couple to sterile right-handed neutrinos via the Dirac mass, and the sterile right-handed neutrinos would couple to their self-conjugate chiral counterpart via the Majorana mass. This situation is described by the following mass-matrix:

$$\mathcal{L} = -\frac{1}{2}(\overline{\nu}_L^c,\ \overline{N_R})\begin{pmatrix} 0 & m_D^T \\ m_D & M \end{pmatrix}\begin{pmatrix} \nu_L^0 \\ N_R^c \end{pmatrix} + \text{h. c.} \tag{1.13}$$

For 3-families of neutrinos m_D and M are 3x3 matrices, while ν_L and N_R are 3-element column vectors.

It is natural to assume that the neutrino Dirac mass m_D is of similar order to those of the other Standard Model particles. It also seems likely that the Majorana mass M is much heavier. In this case, considering for simplicity a single flavour, diagonalising the mass matrix to order $(m_D/M)^2$ gives:

$$\mathcal{M}_\nu = \begin{pmatrix} m_D^2/M & 0 \\ 0 & M \end{pmatrix} \tag{1.14}$$

The mass eigenstates are $m_1 \simeq m_D^2/M$ and $m_2 \simeq M$, giving one very light mass eigenstate and one very heavy one. This is known as the See-Saw model and provides a natural explanation for the lightness of the observed neutrinos. For example with $m_D \sim m_{top}$ and $M \sim 10^{15}\,\text{GeV}$ the light masses $m_1 \simeq 3 \times 10^{-2}\,\text{eV}$, in agreement oscillation experiments. Establishing whether neutrinos are Majorana will require observation of processes that would otherwise be forbidden, such as neutrinoless double beta decay discussed in **??**.

1.3 Fundamental Neutrino Interactions

1.3.1 Gauge Symmetries

Quantum Electro Dynamics (QED) and Quantum ChromoDynamics (QCD) both follow from the symmetry of local gauge invariance, which states that a phase translation of $\Psi \to e^{iq\theta(x)}\Psi$ leaves the Lagrangian invariant. Beginning with QED, there is one scalar electric charge and the symmetry group is U(1). Free charged fermion fields are described by the Lagrangian:

$$\mathcal{L} = i\overline{\Psi}\partial\!\!\!/\Psi - m\overline{\Psi}\Psi \tag{1.15}$$

Imposing local gauge invariance forces one to introduce the notion of a covariant derivative, $\partial_\mu \to D_\mu = \partial_\mu + iqA_\mu$, where A_μ is a gauge field. By imposing that the transformation of this gauge field takes the form $A_\mu \to A_\mu - \partial_\mu\theta(x)$ the Lagrangian remains invariant under local gauge translations. Interpreting this gauge field as the electromagnetic field potential, it becomes apparent that electromagnetic interactions are the physical manifestations of local gauge transitions. Defining the electromagnetic field tensor as $F^{\mu\nu} = \partial^\mu A^\nu - \partial^\nu A^\mu$, a Lagrangian of:

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} \tag{1.16}$$

reproduces Maxwell's equations. The QED Lagrangian for electromagnetically interacting fermions is then given by:

$$\mathcal{L} = \overline{\Psi}(iD\!\!\!\!/ - m)\Psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \tag{1.17}$$

A similar argument leads to the equations of QCD, where the presence of 3 colour charges results in a non-Abelian SU(3) symmetry. The resulting covariant derivative D_μ and gauge fields $G^{\mu a}$ depend on eight generators of the symmetry group, T^a, which describe transitions between the components of the colour multiplets:

$$D_\mu = \partial_\mu + ig_s G^a_\mu T^a \tag{1.18}$$

The colour field tensor takes a similar form to the electromagnetic one but contains eight gauge fields, and its non-Abelian nature results in an additional term describing interactions between these fields.

$$F^k_{\mu\nu} = \partial_\mu G^k_\nu - \partial_\nu G^k_\mu - gf^{klm}G^l_\mu G^m_\nu \tag{1.19}$$

Such self-interactions between gluons leads to quark confinement and the substantially different phenomenologies between QCD and QED, in spite of their otherwise similar equations of motion.

Describing weak interactions through a local gauge symmetry requires new ingredients that will give rise to mass terms for the fermions and massive bosons. These are the introduction of a charged Higgs field and spontaneous symmetry breaking.

Before spontaneous symmetry breaking 3 gauge bosons ($W^{1\mu}$, $W^{2\mu}$, $W^{3\mu}$) mediate SU(2) rotations between left-handed weak doublets, while one further gauge boson (B^{μ}) mediates a U(1) symmetry dubbed 'hypercharge'. The right-handed fermions are SU(2) singlets in line with weak parity violation. By introducing a field charged in the third component of weak isospin and hypercharge (the Higgs field), all these gauge bosons become massive. However, the diagonal components of the generators associated with $W^{3\mu}$ and B^{μ} result in a non-diagonal mass matrix for these bosons. When diagonalized the mass eigenstates are found to describe one massless state (the photon) and one massive one (the Z boson). Combinations of the $W^{1\mu}$ and $W^{2\mu}$ bosons can also be identified with the W^{+} and W^{-} bosons. The fermions gain mass through Yukawa couplings to the Higgs field, and gauge transformations associated with the Higgs field gives rise to the Higgs boson. With appropriate choice of fermion charges one finds this accommodates all necessary features of the electroweak theory - an achievement celebrated by the 2013 Nobel Prize in Physics.

Through electroweak symmetry breaking the the weak force becomes intrinsically linked to electromagnetism, giving the resulting electroweak theory a SU(2)$_L$⊗U(1)$_Y$ symmetry, and a first-generation leptonic Lagrangian of the form:

$$\begin{aligned}\mathcal{L} = i\,\overline{\Psi}\,\not{D}\,\Psi &= i\,\overline{e}\,\not{\partial}\,e + i\,\overline{\nu}\,\not{\partial}\,\nu \\ &- g\sin\theta_W\,\overline{e}\,\gamma^{\mu}\,e\,A_{\mu} \\ &+ \frac{g}{\cos\theta_W}\left(\sin^2\theta_W\,\overline{e_R}\,\gamma^{\mu}e_R - \frac{1}{2}\cos 2\theta_W\,\overline{e_L}\,\gamma^{\mu}\,e_L + \frac{1}{2}\,\overline{\nu_L}\,\gamma^{\mu}\,\nu_L\right)Z^0_{\mu} \\ &+ \frac{g}{\sqrt{2}}\left(\overline{\nu_L}\gamma^{\mu}e_L W^{+}_{\mu} + \overline{e_L}\gamma^{\mu}\nu_L W^{-}_{\mu}\right)\end{aligned} \tag{1.20}$$

1.3.2 Exchange Currents

From Equation 1.20 we can identify the terms involved in neutrino interactions. We have two terms involving the charged and neutral currents j^{CC}_{μ} and j^{NC}_{μ} and the corresponding electroweak exchange bosons $W^{\pm}$ and Z^0.

$$\mathcal{L}^{CC} = -\frac{g}{\sqrt{2}}(j^{CC}_{\mu}W^{\mu} + \mathrm{h.c.}) \tag{1.21}$$

$$\mathcal{L}^{NC} = -\frac{g}{2\cos\theta_W}j^{NC}_{\mu}Z^{\mu} \tag{1.22}$$

The charged current j_μ^{CC} couples only left-handed fields, while the neutral current j_μ^{NC} also couples right-handed fields with a suppression by the weak mixing angle θ_W. This parity violation is expressed through the projection operator $\frac{1}{2}(1-\gamma_5)$, which extracts the left-handed chiral component of the fields. The weak leptonic charged and neutral currents are then:

$$j_\mu^{CC} = \sum_i \bar{l}_\alpha \gamma_\mu (1-\gamma_5) \boldsymbol{V}_{\alpha i}\, \nu_i \tag{1.23}$$

$$j_\mu^{NC} = \sin^2\theta_W \bar{l}\gamma_\mu(1+\gamma_5)l - \frac{1}{2}(\cos 2\theta_W)\bar{l}\gamma_\mu(1-\gamma_5)l + \frac{1}{2}\overline{\nu}_l\gamma_\mu(1-\gamma_5)\nu_l \tag{1.24}$$

The combination $\gamma_\mu(1-\gamma_5)$ involves vector (γ_μ) and axial-vector ($\gamma_5\gamma_\mu$) bilinear covariant terms; the weak charged current is therefore said to be '$V-A$' in nature, while the weak neutral current has a more general '$g_V V - g_A A$' form. For antineutrinos the right-handed projection $\frac{1}{2}(1+\gamma_5)$ results in destructive interference with the transverse part of the vector current, in contrast to constructive interference for neutrinos. This results in a suppression of anti-neutrino interactions relative to those of neutrinos.

The matrix $\boldsymbol{V}_{\alpha i}$ in Equation 1.23 accounts for any mismatch in basis between the weak eigenstates ν_α and mass eigenstates ν_i. In the minimal Standard Model where neutrinos are massless the neutrino states are defined only by their weak interactions, and the matrix $\boldsymbol{V}_{\alpha i}$ can be omitted. The discovery that neutrinos oscillate demonstrated that weak states are not mass eigenstates, and flavour-changing charged-current interactions require a basis translation via the matrix $\boldsymbol{V}_{\alpha i}$ and a sum over mass eigenstates. The matrix $\boldsymbol{V}_{\alpha i}$ is known as the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix and takes the form shown in Equation 1.25.

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \underbrace{\begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix}}_{U_{PMNS}} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{1.25}$$

This may be parameterised in terms of three mixing angles $\theta_{\alpha i}$ and three phases by which CP violation can occur. If neutrinos are Dirac particles just one CP violation phase is physical (the others may be removed through suitable redefinition of the field phases) and is usually denoted the 'Dirac CP phase' δ_{CP}. In terms of these parameters the PMNS matrix may be decomposed into three separate matrices that mix each pair of generations as shown in Equation 1.26

$$U_{PMNS} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & e^{i\delta_{CP}} s_{13} \\ 0 & 1 & 0 \\ -e^{i\delta_{CP}} s_{13} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{1.26}$$

where $c_{\alpha\beta} = \cos\theta_{\alpha\beta}$ and $s_{\alpha\beta} = \sin\theta_{\alpha\beta}$.
For neutral current interactions the flavour composition (whether an eigenstate or superposition of eigenstates) is unchanged, so one can operate in either eigenbasis without any mixing terms and $\boldsymbol{V}_{\alpha i}$ does not appear.

The weak coupling of quarks produces a similar set of currents to those in Equation 1.23 and Equation 1.24. For a quark q with third component of weak isospin t_3 and electric charge e_q, the resulting currents are:

$$j_\mu^{NC} = \sum_q \overline{q}\,\gamma_\mu \left[(t_3 - e_q \sin^2\theta_W)(1-\gamma_5) - e_q \sin^2\theta_W(1+\gamma_5)\right] q \tag{1.27}$$

$$j_\mu^{CC} = \overline{q}_\alpha\,\gamma_\mu(1-\gamma_5)\,q'_\alpha = \overline{q}_\alpha\gamma_\mu(1-\gamma_5)\sum_i \boldsymbol{V}_{\alpha i} q_i = \cos\theta_C \cdot \overline{q}_\alpha\,\gamma_\mu(1-\gamma_5)\,q_\alpha \tag{1.28}$$

As with neutrinos the interaction states are not eigenstates of the free Hamiltonian. The projection required for weak interactions is described by the Cabbibo-Kobayashi-Maskawa (CKM) matrix, which may be decomposed in the same way as the PMNS matrix.

Since neutrino masses are so small the quantum uncertainty in the energy of neutrinos emitted in weak interactions is often larger than the difference in energies of the mass eigenstates. When this is the case neutrinos may produced in a coherent superposition of mass eigenstates, each of which propagates with a different velocity[20]. As the relative contribution from each mass state changes, the corresponding flavour mixture changes too. A subsequent coherent measurement of the neutrino can then project out a different weak state than was initially created - this is the phenomenon of 'neutrino oscillation'. The probability of observing a given state is a function of the neutrino energy, distance of propagation, and the electron density along the path traversed. The probability of changing state (that is, of the neutrino 'oscillating into a different flavour state') is usually factored out of the interaction and calculated separately.
By contrast, the large mass splittings of quarks mean that decoherence times are extremely short[1], and quarks can generally be assumed to occupy mass eigenstates. The corresponding projection is typically incorporated into the weak coupling of the down-type quarks. Although equivalent in origin, this distinguishes quark *mixing* from neutrino *oscillation.* The degree of mixing is also substantially lower in the quark sector than for neutrinos (Figure 1.1), and it is often sufficient to take only the dominant contribution from the diagonal elements. Dropping Cabbibo suppressed and doubly suppressed terms reduces the sum over mass eigenstates to the cosine of the Cabbibo angle $\theta_C = \theta_{12}$, resulting in the last equality of Equation 1.28.

---

[1] there are some exceptions, such as the case of Kaon mixing

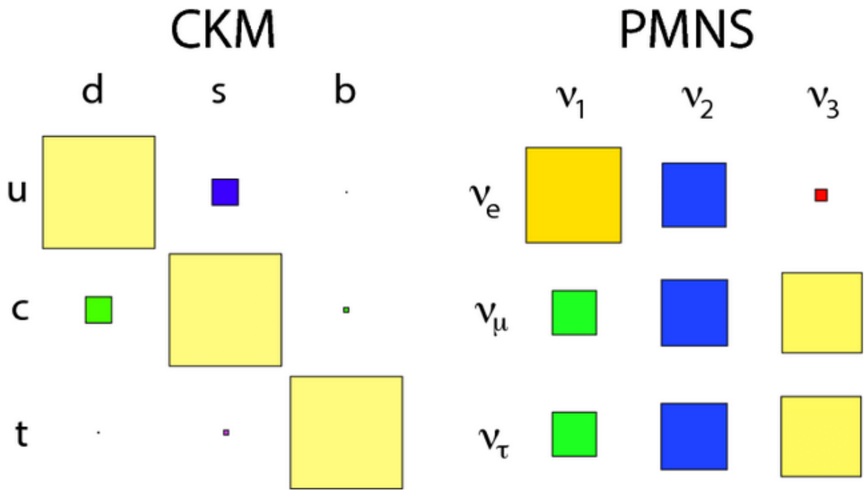

Figure 1.1: Relative degrees of mixing between generations of quarks and neutrinos, represented by the elements of the CKM and PMNS matrices respectively. Figure from [21].

1.3.3 Quasi-Elastic Scattering

Since there are no free quarks, the simplest hadronic interaction we can consider is the weak isospin transition of a free nucleon, $\nu + n \rightarrow l^- + p^+$ and $\overline{\nu} + p \rightarrow l^+ + n$ (Figure 1.2).

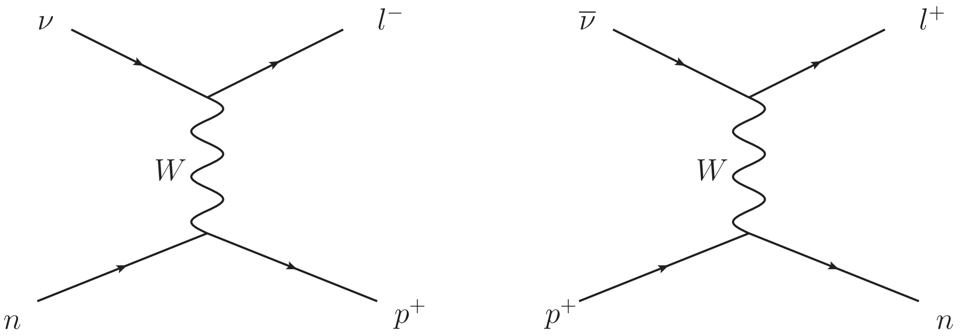

Figure 1.2: Quasi-elastic neutrino scattering. For low energy $\overline{\nu_e}$ this interaction is often called 'inverse beta decay'.

In order to account for the substructure of the nucleon, the hadronic currents in Equation 1.27 and Equation 1.28 must be modified.

By applying the required symmetries of the Lagrangian it is possible to write the hadronic current generally in terms of six 'form factors' with specific Lorentz transformation properties, each dependant only on the squared momentum transfer Q^2 and the particle masses. Explicitly splitting the charged current into its '$V - A$' parts, the vector and axial-vector components may be written in terms of the nucleon spinor u, nucleon mass M and the form factors F_i as:

$$j^h_\mu = \cos\theta_C(V_\mu - A_\mu) \tag{1.29}$$

$$V_\mu = \bar{u}\left[\gamma_\mu F^V_1(Q^2) + \frac{i}{2M}\sigma_{\mu\nu}q^\nu F^V_2(Q^2) + \frac{q_\mu}{M}F^S(Q^2)\right]u \tag{1.30}$$

$$A_\mu = -\bar{u}\left[\gamma_\mu\gamma_5 F^A(Q^2) + \frac{q_\mu}{M}\gamma_5 F^P(Q^2) + \frac{i}{2M}\sigma_{\mu\nu}q^\nu\gamma_5 F^T(Q^2)\right]u \tag{1.31}$$

Two of these form factors, the tensor term F^T and scalar term F^S, are zero in the absence of second class currents and can safely be neglected[22].

Insight into the vector form factors, F^1_V and F^2_V may be gained from the scattering of charged leptons on nucleons. Electromagnetic interactions do not violate parity, so contain only a vector current term usually parameterised in terms of the so-called Dirac and Pauli form factors of the proton and neutron, representing the deviation from a point charge and point magnetic moment respectively. In the Conserved Vector Current (CVC) hypothesis the weak charged current is decomposed into its vector and axial-vector parts, and the electromagnetic current is decomposed into isovector and isoscalar parts. It is then possible to identify 3 components of a conserved current in isospace; two parts make up the weak vector current and the remaining part is the electromagnetic isovector term. By proposing that these are parts of the same current, the weak vector currents can be related to the electromagnetic isovector current, being related through a rotation in isospace. As a result the corresponding form factors must be the same, and so the weak form factors can be related to the Pauli and Dirac form factors for the proton and neutron. The resulting relations are:

$$\begin{aligned} F^V_1(Q^2) &= F^p_{Dirac}(Q^2) - F^n_{Dirac}(Q^2) \\ F^V_2(Q^2) &= F^p_{Pauli}(Q^2) - F^n_{Pauli}(Q^2) \end{aligned} \tag{1.32}$$

These equalities hold in the limit of degenerate quark masses, since they are based on isospin symmetry. The electromagnetic form factors are well established through electron scattering experiments[23].

A similar current conservation law can be found for the axial-vector components of the weak charged and neutral currents. At the quark level the divergence of the axial-vector current is related to the difference in quark masses, such that the axial-vector current is conserved for degenerate quarks. Extending this to the case of nucleons it is possible to demonstrate that the divergence is instead proportional to the pion mass[24], so conservation holds in the looser limit of $m_\pi \to 0$ for nucleon interactions. This is known as the 'Partially Conserved Axial-vector Current', and allows one to relate the pseudoscalar form factor F_P to the axial-vector form factor F_A[25].

$$F_P(Q^2) = \frac{2M^2}{Q^2 + m_\pi^2} F_A(Q^2) \tag{1.33}$$

This leaves one remaining form factor that must be extracted from weak interactions. A significant amount of early and current work gives the axial-vector form factor a dipole form, parameterised by the axial mass M_A:

$$F_A(Q^2) = g_A \frac{1}{\left(1 + \frac{Q^2}{M_A^2}\right)^2} \tag{1.34}$$

More recently it has become clear that fits of experimental data to this basic parameterisation result in a significant under-estimation of systematic errors. A more general method is the 'z-expansion'[26], which provides improved (i.e. larger) error estimates, but has no predictive power and does not produce a better fit than the dipole ansatz[27]. Better predictive models of F_A may come from Lattice QCD, where work is ongoing[28].
In the limit of zero momentum transfer, $Q^2 \to 0$, the form-factors reduce to constants:

$$\begin{aligned} F_1^V(0) &\to 1 \\ F_2^V(0) &\to \frac{\mu_p - \mu_n}{\mu_N} - 1 \simeq 3.706 \\ F_A(0) &\to -g_A \end{aligned} \tag{1.35}$$

where μ_n and μ_p are the nucleon magnetic form-factors, μ_N is the nuclear magneton, and g_A is the axial coupling constant. The interaction in this limit is typically referred to as 'inverse beta decay'.

1.3.4 Resonant Pion Production (RES)

As well as the simple isospin transitions described above, in the presence of a nucleus it is also possible for neutrinos to induce the production of mesons, most often pions. This becomes kinematically accessible for neutrino energies of ~0.5 GeV and above, and by ~1 GeV this mechanism becomes the dominant contributor to the total cross-section. The process proceeds through the resonant excitation of a heavy baryon followed by its subsequent decay into a pion and nucleon (Figure 1.3).

The heavy baryon is predominantly the (spin-3/2) $\Delta(1232)$, although numerous other resonances, both spin-1/2 and spin-3/2, are accessible for πN invariant masses below 2 GeV.

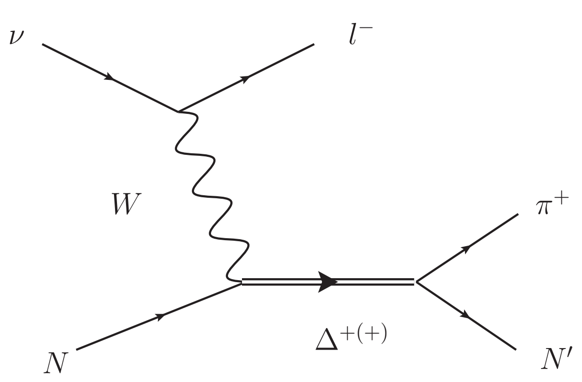

Figure 1.3: Resonant single pion production. N is the incoming nucleon and N' the outgoing nucleon.

For spin-1/2 resonances the current may be factorised similarly to Equation 1.30 and Equation 1.31, with the equivalent forms being[29]:

$$V^\mu = \frac{F_1}{(2M_N)^2}(Q^2\gamma^\mu + \not{q}q^\nu) + \frac{F_2}{2M_N} i\sigma^{\mu\alpha}q_\alpha \tag{1.36}$$

$$-A^\mu = F_A\gamma^\mu\gamma_5 + \frac{F_P}{M_N}q^\mu\gamma^5 \tag{1.37}$$

It should be noted that the form factors above are *not* the same as those in Equation 1.30 and Equation 1.31. Likewise, while resonant production can occur through CC, NC and EM channels, the form factors in each case will again be different, although it is possible to relate them to each other. Finally, while the total current has a form $J^\mu = (V^\mu - A^\mu)$ for excitation of positive parity resonances, an additional factor arises for production through negative parity resonances, giving $J^\mu = (V^\mu - A^\mu)\gamma^5$.

For spin-3/2 resonances the associated baryon states are vector spinors, and the currents may be expressed in the form:

$$V^{\alpha\mu} = \frac{F_1^V}{M_N}(g^{\alpha\mu}\not{q} - q^\alpha\gamma^\mu) + \frac{F_2^V}{M_N^2}(g^{\alpha\mu}q\cdot p' - q^\alpha p'^\mu) + \frac{F_3^V}{M_N^2}(g^{\alpha\mu}q\cdot p - q^\alpha p^\mu) + g^{\alpha\mu}F_4^V \tag{1.38}$$

$$-A^{\alpha\mu} = \left[\frac{F_1^A}{M_N}(g^{\alpha\mu}\not{q} - q^\alpha\gamma^\mu) + \frac{F_2^A}{M_N^2}(g^{\alpha\mu}q\cdot p' - q^\alpha q'^\mu) + F_3^A g^{\alpha\mu} + \frac{F_3^A}{M_N^2}q^\alpha q^\mu\right]\gamma_5 \tag{1.39}$$

The same remarks for the form factors apply, but this time the currents are of the form $(V^{\alpha\mu} - A^{\alpha\mu})$ for *negative* parity resonances, and $J^{\alpha\mu} = (V^{\alpha\mu} - A^{\alpha\mu})\gamma^5$ for *positive* parity resonances.

As with the QE case, the CVC hypothesis is used to relate the charged-current vector form factors to those for electromagnetic resonant production, which can be measured from electron scattering helicity amplitudes. The PCAC hypothesis relates the pseudo-scalar form factor to the axial-vector one, for which a dipole form is assumed. The coupling constant $F_A(0)$ can be extracted from pion scattering cross-section measurements, leaving an axial mass parameter to be determined from data[30].

As well as resonant excitation it is also possible for non-resonant 'background' processes to generate the intermediate heavy baryon. For isospin-3/2 baryons this background is small and often neglected, being below the level of other uncertainties. In the case of the isospin-1/2 excitations, however, the non-resonant contribution is of similar size to the resonant one and must be accounted for to achieve good agreement with data. This background may be taken from the deficit between prediction and data in spin-1/2 production channels.

1.3.5 Deep Inelastic Scattering (DIS)

As the neutrino energy increases it becomes possible to generate multiple pions and, eventually, hadronic cascades (Figure 1.4). This realm is known as deep inelastic scattering. This begins to become relevant for energies above 2 GeV, becoming dominant at energies above 5 GeV.

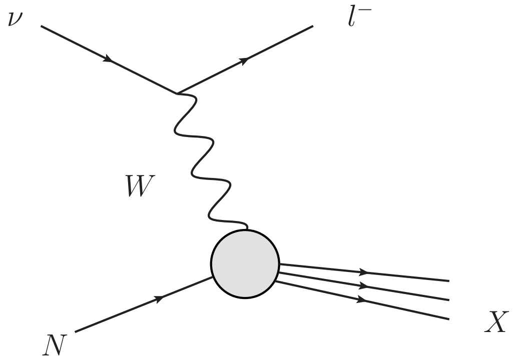

Figure 1.4: Deep Inelastic Scattering. N is the incoming nucleon and X an outgoing hadronic cascade.

At these energies the neutrino is able to probe the quark substructure of the nucleon, so the form factors are related to the nucleon parton distribution functions (PDFs) describing the probability of finding a target quark with a given flavour and momentum. These PDFs are parameterised in terms of the interaction inelasticity y, describing the fractional energy loss of the probe, and the Björken scaling variable $x \equiv Q^2/2MEy$. These PDFs have been extensively studied in lepton-hadron scattering experiments, and the cross-sections and form factors may be found in the literature[31][32].

At the lower end of this scale there is some overlap with resonant pion production, and the transition between models can result in discontinuities that are undesirable in Monte Carlo generation. To mitigate this generators often include a sample of DIS events in the the non-resonant part of single pion production. A particularly suitable DIS parameterisation in this case is the Bodek-Yang parameterisation, which enables a coherent description of resonant and DIS events together[33][34]. DIS scattering can be considered approximately as elastic scattering from free quarks, a situation that results in the 'scaling' property whereby the structure functions are independent of Q^2. The asymptotically-free nature of QCD gives rise to corrections to this scaling, which may be fit with some form factor. Bodek and Yang use a modified scaling function, which allowed 12 parameters to fit resonances and 8 parameters to account for low W behaviour. These parameters were obtained from electron-proton scattering data, resulting in a set of PDFs that are able to model both DIS and resonant single-pion production.

1.3.6 Amplitudes and Cross-Sections

In order to derive scattering amplitudes and cross-sections we may use perturbation theory to write the transition between two free states in terms of the scattering matrix element. Explicitly denoting the current that connects an initial state $|i(k)\rangle$ with momentum k to final state $|f(k')\rangle$ with momentum k' as $\langle\, i(k)\,|\, j_\mu \,|\, f(k')\,\rangle$, the scattering amplitude for the interaction $\nu(k) + n(p) \rightarrow l(k') + p(p')$ may be written as

$$\mathcal{M} = \left(\frac{g}{2\sqrt{2}}\right)^2 \langle\, l(k')\,|\, j_l^\mu \,|\, \nu(k)\,\rangle \; \frac{i}{q^2 - M_W^2}\left(-g^{\mu\nu} + \frac{q^\mu q^\nu}{M_W^2}\right) \langle\, p(p')\,|\, j_h^\mu \,|\, n(p)\,\rangle \tag{1.40}$$

For low momentum transfers we can simplify the propagator term

$$\frac{i}{q^2 - M_W^2}\left(-g^{\mu\nu} + \frac{q^\mu q^\nu}{M_W^2}\right) \rightarrow \frac{ig^{\mu\nu}}{M_W^2} \tag{1.41}$$

to give

$$\mathcal{M} = i\frac{g^2}{8M_W^2} \; \langle\, l(k')\,|\, j_l^\mu \,|\, \nu(k)\,\rangle \langle\, p(p')\,|\, j_\mu^h \,|\, n(p)\,\rangle \tag{1.42}$$

The coefficient is often then written in terms of the Fermi coupling, $g^2/8M_W^2 = G_F/\sqrt{2}$. Transition probabilities are the square of the amplitude:

$$|\mathcal{M}|^2 = \frac{G_F^2}{2} \; \langle\, p(p')\,|\, j_h^\mu \,|\, n(p)\,\rangle \langle\, n(p)\,|\, j_h^\nu \,|\, p(p')\,\rangle \; \langle\, l(k')\,|\, j_{\mu,l} \,|\, \nu(k)\,\rangle \langle\, \nu(k)\,|\, j_{\nu,l} \,|\, l(k')\,\rangle \tag{1.43}$$

The differential interaction cross-section is obtained from the transition probabilities after accounting for the phase space and flux factors:

$$d\sigma = \frac{1}{4(k \cdot p)}(2\pi)^4 \delta^4(p - k + p' - k')\frac{d^3k'}{(2\pi)^3 2E'_k}\frac{d^3p'}{(2\pi)^3 2E'_p}|\overline{\mathcal{M}}|^2 \tag{1.44}$$

Converting the Lorentz-invariant phase space integration to the square momentum transfer q^2 and integrating over outgoing angles gives:

$$(2\pi)^4 \delta^4(p - k + p' - k')\frac{d^3k'}{(2\pi)^3 2E'_k}\frac{d^3p'}{(2\pi)^3 2E'_p} = \frac{dq^2}{16\pi M E_k} \tag{1.45}$$

The matrix elements are usually encapsulated in terms of the leptonic tensor $\mathcal{L}_{\mu\nu}$ and hadronic tensor $\mathcal{W}_{\mu\nu}$:

$$\mathcal{L}^{\mu\nu} = \sum_{\text{spins}} \langle\, l(k') \,|\, j_l^{\mu} \,|\, \nu(k) \,\rangle \langle\, \nu(k) \,|\, j_l^{\nu} \,|\, l(k') \,\rangle \tag{1.46}$$

$$\mathcal{W}^{\mu\nu} = \overline{\sum_{\text{spins}}} \int \langle\, p(p') \,|\, j_h^{\mu} \,|\, n(p) \,\rangle \langle\, n(p) \,|\, j_h^{\nu} \,|\, p(p') \,\rangle \tag{1.47}$$

where the line over the sum indicates an average over initial spins and the sums runs over final spins. The differential cross-section with respect to Q^2 can then be written:

$$\frac{d\sigma}{dQ^2} = \frac{1}{64\pi E_\nu^2 M^2}\frac{G_F^2}{2}\frac{\cos\theta_C^2}{2}\mathcal{L}^{\mu\nu}\mathcal{W}_{\mu\nu} \tag{1.48}$$

Since the states $|l\rangle$, $|\nu\rangle$ are simple Dirac spinors the leptonic tensor may be calculated explicitly using completeness relations. The result, neglecting the lepton mass, is:

$$\begin{aligned}
\mathcal{L}_{\nu\mu} &= \sum_{\text{spins}} [\overline{\nu}_l \gamma_\mu (1 - \gamma_5) l][\bar{l}\gamma_\mu (1 - \gamma_5)\nu] \\
&= Tr[\not{k}'\gamma_\mu (1 - \gamma_5)\not{k}\gamma_\nu (1 - \gamma_5)] \\
&= 2\left(k_\mu k'_\nu + k'_\mu k_\nu - g_{\mu\nu} k_\alpha k^\alpha - i\epsilon_{\mu\nu\alpha\beta} k^\alpha k'^\beta\right)
\end{aligned} \tag{1.49}$$

where $g_{\mu\nu}$ is the metric tensor and $\epsilon_{\mu\nu\alpha\beta}$ is the fully antisymmetric Levi-Civita tensor.

In the case of the hadronic tensor there are a few complications. The main issue is the determination of the form factors, which depend on the specific interaction mode and channel. For resonant pion production the finite width of the resonance also generalises the delta function in Equation 1.44 to a Breit-Wigner distribution:

$$\delta^4(p - p' - q) \rightarrow \frac{p'}{\pi} \frac{\Gamma(p')}{(p'^2 - M_R^2)^2 + p'^2\Gamma^2(p')} \tag{1.50}$$

where M_R is the mass of the resonant baryon. Finally when considering spin-3/2 resonances the completeness relations must be calculated using the Rarita-Schwinger formalism.

1.3.7 Neutrino Scattering on Charged Leptons

The simplest case we can consider is neutrino-lepton scattering. Following the procedure above with two leptonic currents, the total cross-section for $\nu_\mu + e^- \rightarrow \nu_e + \mu^-$ is:

$$\sigma \simeq \frac{2m_e G_F^2 E_\nu}{\pi} \tag{1.51}$$

Other neutrino-lepton scattering cross-sections can be written relative to Equation 1.51 with various coefficients involving the weak mixing angles and isospins.

A useful property of these interactions is that the outgoing charged lepton closely follows the direction of the incoming neutrino; $E_e\theta_e^2 < 2m_e$. This makes neutrino-lepton scattering a very useful channel for locating an unknown neutrino source. This may be applied to $\nu_e - e$ scattering to identify the location of a supernovae, or the direction of a reactor in non-proliferation detection experiments, for example. Unfortunately the dependence on m_e means the cross-section for these interactions is incredibly small; at $E_\nu \sim 1\,\text{GeV}$ interactions on electrons are $\sim 10,000$ times less likely than those on nucleons.

1.3.8 Quasi-Elastic Scattering on Nucleons

The cross-sections that follow from Equation 1.30 - Equation 1.34 were first worked out by Llewellyn-Smith[35], who expressed the resulting charged current quasi-elastic (CCQE) cross-section as follows:

$$\frac{d\sigma}{dQ^2} = \frac{G_F^2 M^2 |V_{ud}|^2}{8\pi E_\nu^2} \left[A \pm \frac{(s-u)}{M^2} B + \frac{(s-u)^2}{M^4} C\right] \tag{1.52}$$

$$\begin{aligned} A(q^2) = \frac{m^2 - q^2}{M^2} \Bigg[&(1+\eta)F_A^2 - (1-\eta)F_1^2 + \eta(1-\eta)F_2^2 + 4\eta F_1 F_2 \\ &- \frac{m^2}{4M^2}\left((F_1 + F_2)^2 + (F_A + 2F_P)^2 - (\eta + 4)\,F_P^2\right)\Bigg] \end{aligned} \tag{1.53}$$

$$B(q^2) = \eta F_A(F_1 + F_2) \tag{1.54}$$

$$C(q^2) = \frac{1}{4}(F_A^2 + F_1^2 + \eta F_2^2) \tag{1.55}$$

The $(-)+$ refers to (anti)neutrinos, $\eta \equiv -q^2/4M^2$, with M the nucleon mass, s and u are the usual Mandlestam variables and m is the lepton mass.
In the case of neutral current elastic scattering the pseudoscalar term is suppressed by the outgoing lepton mass so that the cross-section is ~10-20% that of CCQE.

CCQE is the golden channel for neutrino interactions as it allows one to infer the neutrino energy and momentum transfer from just the detected lepton kinematics, at least while neglecting nuclear dynamics and final state interactions (these will be discussed later - the Llewellyn-Smith formula in Equation 1.52 describes the scattering of neutrinos on *free* nucleons). The derived neutrino energy and momentum transfer are given by:

$$E_\nu = \frac{2(M_n - E_B)E_l - (E_B^2 - 2M_nE_B + m_l^2 + \Delta M^2)}{2\big[M_n - E_B - E_l + |\vec{p}_l|cos\theta\big]} \tag{1.56}$$

$$Q^2 = 2E_\nu(E_l - |\vec{p}_l|\cos\theta) - m_l^2 \tag{1.57}$$

Here M_n, M_p and m_l are the neutron, proton and lepton masses, with $\Delta M^2 = M_n^2 - M_p^2$, E_l and θ are the energy and emission angle of the lepton, with $|\vec{p}_l| = \sqrt{E_l^2 - m_l^2}$, and E_B is the binding energy.
For this simple two-body scattering the outgoing lepton angle is directly related to the energy transfer, and if the outgoing nucleon is to be on-shell then the energy and momentum transfers are also related. The mapping from neutrino energy and lepton angle (E_ν, θ), to energy and momentum transfers, (ν, Q^2), is therefore a 1:1 relation. A plot of $\frac{d\sigma}{dQ^2}$ would be a simple delta-Dirac function. As we will see later, nuclear dynamics smear this delta into a broadened distribution known as the 'quasi-elastic peak'. This is shown in Figure 1.5 ('QEL', blue curve), along with the cross-section contributions from resonant pion production and deep inelastic scattering, discussed in the next sections.

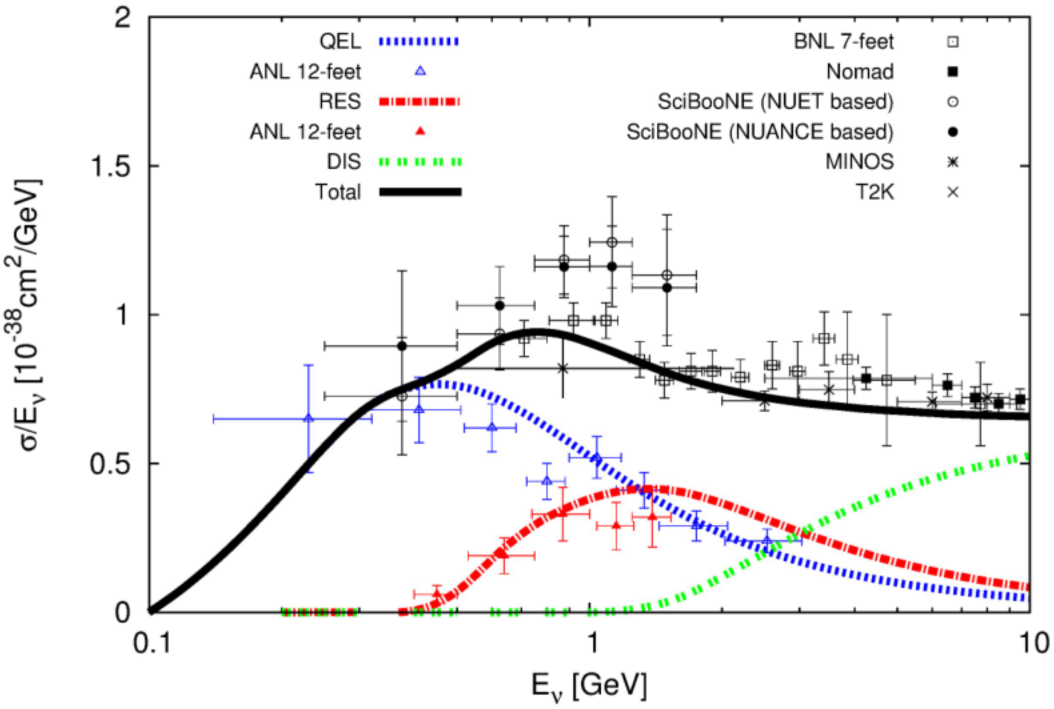

Figure 1.5: Total neutrino interaction cross-section, along with estimated contributions from quasi-elastic scattering (QEL), resonant pion production (RES), and deep inelastic scattering (DIS). Points show data from various experiments as indicated. Figure from [36].

1.3.8.1 Resonant Pion Production

Due to the complexity and variety of approaches used in modelling pion production processes, no standard equivalent to the Llewellyn-Smith equation exists for pion production cross-sections. Nonetheless methods akin to that described in subsection 1.3.4, with form factors fit to electron scattering and pion production data, are able to provide a reasonably accurate description for the $W_{\pi N} < 1.4\,\mathrm{GeV}$ region[37].

The most widely used model for calculating cross-sections, however, is that of Rein and Seghal[38], which takes a slightly different approach. Originally developed in 1980, Rein and Seghal used the simplified quark harmonic oscillator model of Feynman, Kislinger and Ravndal (the FKR model) to theoretically calculate matrix elements for 18 baryon resonances at invariant masses below 2 GeV. These were then combined, including interference terms, with an empirical non-resonant background component. At the time the model did well at reproducing the available data, but as experiments became more precise and a wider range of invariant masses was explored the model has been shown to be lacking. It substantially under-estimates pion production in electron-scattering experiments[39], and does not agree well with more accurate neutrino scattering models[40], under-estimating the Δ resonance peak and over-predicting at higher invariant masses (Figure 1.6). It also provides only a cross-section prediction, but no information on the produced pion kinematics. In spite of these shortcomings its ease of implementation means that the Rein-Seghal model is still the default method of calculation in many Monte Carlo generators.

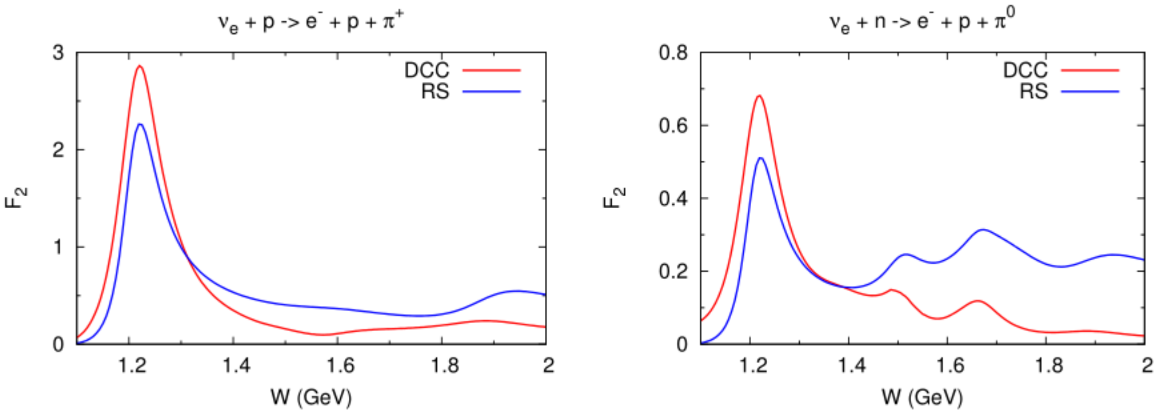

Figure 1.6: Comparison of cross-sections for forward scattering ($\theta = 0$) pion production events, between the Rein Seghal model (RS) and a more realistic Dynamical Coupled Channel model (DCC)[41]. Plots show $W\frac{d\sigma}{dEd\Omega}$ where W is the energy of the pion-nucelus system in its rest frame.

Neutrino energy reconstruction in pion production events is generally done through calorimetric means. The total energy of all outgoing products is summed up and equated to the kinetic energy provided by the neutrino. Such methods are well-suited to tracking detectors with a low energy threshold, but Cherenkov detectors must rely on the observable kinematics of high-energy charged particles to reconstruct the event. If one is able to determine the kinetic energy of the pion the neutrino energy can be estimated from:

$$E_\nu = \frac{m_\mu^2 + m_\pi^2 - 2M(E_\mu + E_\pi) + 2p_\mu \cdot p_\pi}{2(E_\mu + E_\pi - |\vec{p}_\mu|cos\theta_{\mu-\nu} - |\vec{p}_\pi|cos\theta_{\nu-\pi} - M)} \tag{1.58}$$

where the target nucleon is taken at rest. For muon neutrino interactions a common difficulty in accomplishing this is the correct identification of the pion and lepton in the event; their similar masses result in very similar light distributions. Using the incorrect identity, however, results in a significant shift in the reconstructed energy. MiniBooNE demonstrated a likelihood based identification using a technique based on events in which the charged pion underwent hadronic scattering[42]. The pattern of photon hits were fit with a likelihood function that accommodated a scattering point, resulting in two distinct sections of emission. This represented the first time charged pions had been tracked and their kinematics measured in a Cherenkov detector.

1.4 Neutrino-Nucleus Interactions

Early neutrino detection experiments used light nuclei such as hydrogen or deuterium and were reasonably well approximated by simple adaptations of the free nucleon model. Unfortunately the statistics possible with these targets are very limited owing to the small cross-sections of light nuclei, and later experiments moved to the use of heavier nuclei (carbon, oxygen, iron etc) to increase cross-sections and so reaction rates. Although this alleviates the issue of low statistics, it opens a Pandora's box of nuclear complications which continue to be a significant challenge to modern experiments.

1.4.1 The Plane Wave Impulse Approximation

The simplest assumption we can make is that the nucleons are independent and quasi-free, and that the interaction occurs with just one nucleon (the impulse approximation) with no further effects arising from traversal within the nuclear medium (the plane-wave approximation). In this case the hadronic tensor for the nuclear interaction is treated as an incoherent sum over those of the potential nucleon interactions:

$$W_A^{\mu\nu} = \int d^3r \, \rho(r) \int \frac{d^3p \; dE}{E_p \, E_{p'}} P(\vec{p}, E) \; W_i^{\mu\nu}(\tilde{q}, p) \tag{1.59}$$

$W_A^{\mu\nu}$ is the nucleus hadronic tensor, $W_i^{\mu\nu}$ is the nucleon hadronic tensor, and the integrals run over the spatial, momentum and energy distributions of the nucleons. Note the momentum transfer $\tilde{q}$ is modified to account for the fraction of energy lost to place the nucleon on-shell. The 'Spectral Function' $P(\vec{p}, E)$ describes the energy-momentum distribution of target nucleons.

1.4.2 Spectral Functions

The spectral function needs to account for the overlap integral between the wavefunctions representing the nuclear ground state and all possible final states associated with the observed interaction. That is, if $|\Psi_0^A\rangle$ represents the ground state of the initial nucleus with energy E_0, $|\Psi_n^{A-1}\rangle$ represents the nth eigenstate of the final state nucleus with energy E_n, and $a_{\vec{p}}$ represents the annihilation operator of a nucleon with momentum $\vec{p}$, then the spectral function corresponds to:

$$P(\vec{p}, E) = \sum_n \left| \langle \Psi_n^{A-1} | a_{\vec{p}} | \Psi_0^A \rangle \right|^2 \delta(E + E_0 - E_n) \tag{1.60}$$

A full calculation of this spectrum is a complex many-body problem. For sufficiently small nuclei the spectral function can be calculated using Greens Function Monte Carlo (GFMC), or

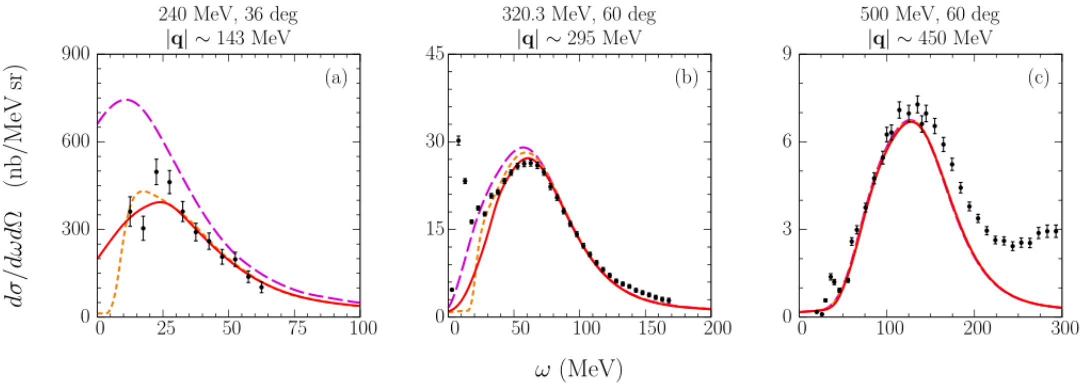

Figure 1.7: The effect of Pauli blocking is a suppression of the cross-section at low energy transfers. Dashed lines use a Global Fermi Gas without Pauli blocking, solid lines use a Global Fermi Gas including Pauli Blocking, and dotted lines a Local Fermi Gas including Pauli blocking. Figure from [46]

for some larger nuclei (including oxygen [43]) Variational Monte Carlo may obtain solutions to a very good approximation[44]. These simulations are computationally expensive and require a defined set of kinematics, so while they are a good point of reference most analyses rely on more analytical approaches.

1.4.2.1 Fermi Gas Model

The simplest spectral function is the Fermi Gas model initially employed by Smith and Moniz[45]. The model assumes a uniform distribution of nucleon momenta up to the Fermi surface ('FG' in Figure 1.10 right) and accounts for a mean binding energy B. This can be expressed as:

$$P(\vec{p}, E) = \frac{3A}{4p_F^3}\Theta(p_F(\rho) - |\vec{p}|)\,\delta(E + \sqrt{M^2 + p^2} - B)\,\Theta(p_F - |\vec{k} - \vec{q}|) \tag{1.61}$$

Here Θ is the Heaviside step function. The Fermi momentum p_F may either be taken as a constant (Global Fermi Gas), or as a function of the local nuclear density ρ (Local Fermi Gas), and may be calculated using the Schrödinger equation (Fermi Gas Model) or the Dirac equation (Relativistic Fermi Gas). The final factor, $\Theta(p_F - |\vec{k} - \vec{q}|)$, accounts for the occupancy of nucleon states following the interaction, such that transitions to states which are already occupied are forbidden - this is known as Pauli Blocking, and becomes important at momentum transfers below $\mathcal{O}(p_F)$ (Figure 1.7).

1.4.3 Nucleon Correlations

The assumption that nucleons behave independently within a nucleus is an approximation that breaks down both at high and low momentum transfers, owing to two different forms of nucleon

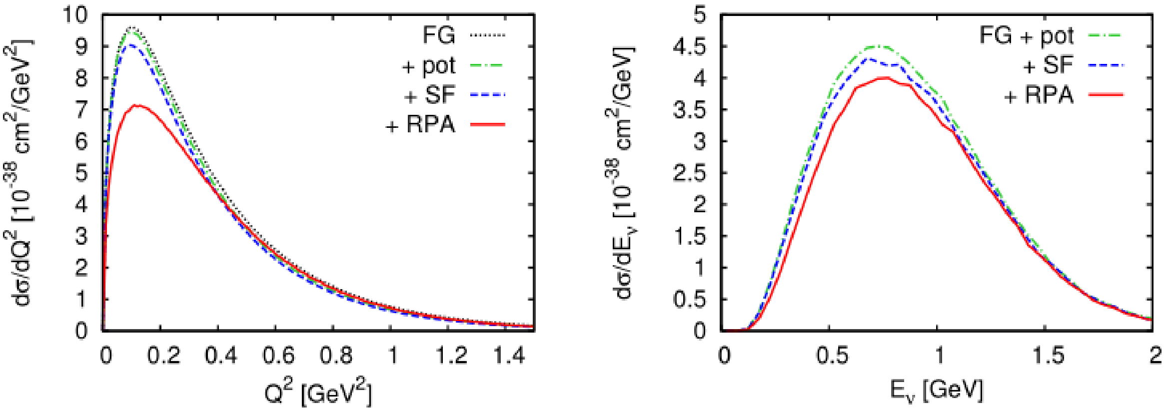

Figure 1.8: The effect of long range correlations (RPA) is a suppression of the cross-section at low energy transfers. These figures show the impact on the true CCQE cross-section averaged over the MiniBooNE flux. 'FG' represents results obtained with a local Fermi Gas model, '+pot' indicates the addition of a mean field contribution to the nucleon initial and final state momenta distributions, 'SF' indicates the further addition of short-range correlations to the outgoing nucleon states, and finally 'RPA' indicates the addition of long range correlations. Figure from [51]

correlations.

1.4.3.1 Long Range Correlations

Strong interactions between nucleons generate a polarization of the nuclear medium which modifies the weak charge of nucleons relative to those in vacuum[47]. Gell-Man and Brueckner originally demonstrated that this effect arises when considering interactions of the propagator with the medium. Its effect can be calculated by replacing the bare propagator with the sum of a series of increasingly long chains of interactions between particle-hole states in the medium[48]. Owing to the context of its original development, this calculation is known as the 'Random Phase Approximation' (RPA). The result of replacing the weak exchange boson with such an expansion is referred to as a 'dressed propagator'. Its impact on neutrino cross-sections has been studied by Nieves et. al.[49][50]. As shown in Figure 1.8 this produces a suppression of the cross-section that is most prominent at low momentum transfers. As the momentum transfer increases the interaction probes smaller distances, and the sensitivity to these correlations diminishes.

1.4.3.2 Short Range Correlations

The interplay between different components of the nucleon-nucleon hadronic tensor is such that the potential between two nucleons is repulsive at very short distances but attractive at distances $\gtrsim$1 fm. This impacts the joint probability of finding two nucleons at a given separation,

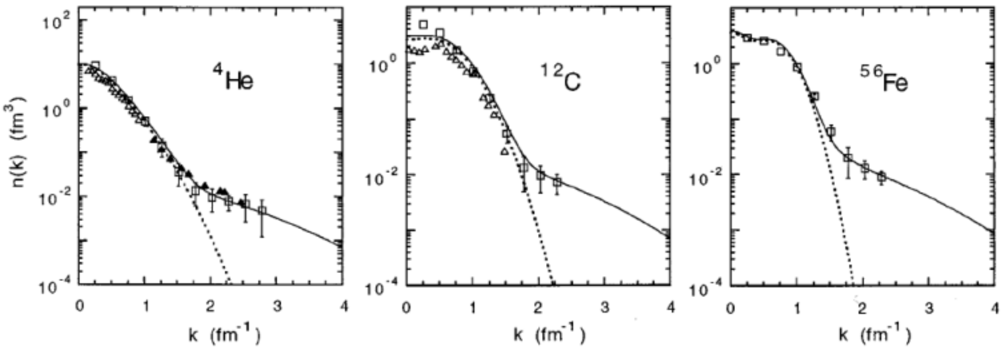

Figure 1.9: The dotted lines indicate the contribution to the momentum distribution arising from the mean field component. The additional high energy tail in the total distributions (solid line) arise from short range two-body interactions, and are seen to be universal for all nuclei. Figure modified from [55]

resulting in a $\mathcal{O}(20\%)$ probability of nucleons forming pairs (typically proton-neutron pairs[52]). These pairs can have a large relative momentum with only a small centre of mass momentum, introducing a high tail into the nucleon momentum distribution function that extends well beyond the Fermi momentum. Since it arises from two-body interactions this tail is independent of the size of the nucleus, as seen in (Figure 1.9). One expects that when one member of a correlated pair is ejected its partner, left with an unbalanced momentum well above the Fermi level, will also often be ejected. Experimental efforts to identify such events and explicitly measure the contribution from such pairs to the spectral function are ongoing[53][54].

1.4.4 Benhar's Spectral Function

Given the universality of short range correlations, one way to simplify the calculation of an analytical spectral function is to decompose it into the sum of two parts; a mean-field part and a short-range correlation part[56]. The first component is calculated by writing the initial and final states of the nucleus as a Slater determinant of shell states within a mean field. The overlap integral then collapses into a sum of single-particle orbitals, each normalised by the occupancy of that orbital. These normalizations may be obtained from nucleon knock-out experiments. The correlation component is in general difficult to calculate for an arbitrary nucleus, but thanks to its universality, it may be taken from the simplified case of infinite, uniform density nuclear matter. This is then integrated over the density profile of the target nucleus under the Local Density Approximation[56] to obtain an estimate for the contribution in the case of the target nucleus. The results of this calculation for oxygen carried out by Benhar[57] demonstrated good agreement with variational Monte Carlo results, as shown in Figure 1.10.

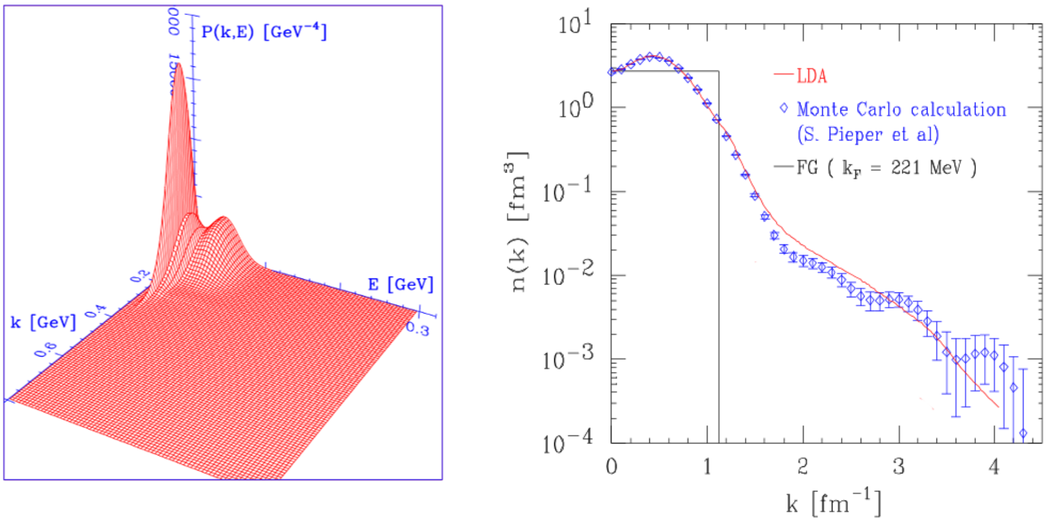

Figure 1.10: Left: Spectral function for ^{16}O calculated with the method of subsection 1.4.4. Right: resulting energy-integrated momentum distribution (LDA), along with those of the Fermi Gas (FG) model and Variational Monte Carlo simulation for comparison. Figures from [57]

1.4.5 Multi-Nucleon Interactions

Continuing down the rabbit hole, yet further complications arise when we reconsider the validity of the impulse approximation. In dilute, weakly interacting matter the approximation that the exchange boson interacts with just one nucleon is well justified. In a nuclear environment, where the density is high and particles are strongly interacting, this approximation breaks down and one must consider more complex many-body interactions in which the exchange boson couples to multiple nucleons at once. This was famously demonstrated by the 'MiniBooNE M_A Puzzle', in which the experiment measured a significantly larger axial mass ($\sim$1.35 GeV) than the previous world average ($\sim$1.03 MeV). This measurement was inferred from the differential charged-current cross-section measured in quasi-elastic events[58]. An analysis by Martini et. al[59] demonstrated that 2-particle 2-hole (2p2h) processes, in which the transferred energy and momentum is shared between two nucleons, could explain the observed discrepancy without an inflated axial mass. The results of predictions with and without these 2p2h processes are shown in Figure 1.11.

The authors utilised the optical theorem, which expresses the differential cross-section in

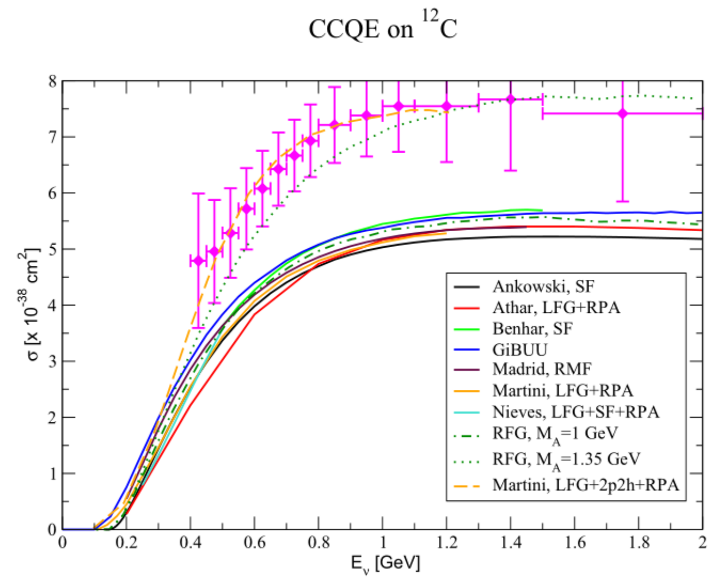

Figure 1.11: Charged-current quasi-elastic cross-section per nucleon as a function of neutrino energy. The lines show various nuclear models, but only the Martini model including 2p2h contributions matches the data from MiniBooNE. Figure from [59]

terms of the exchange boson self-energy tensor $\Pi^{\nu\mu}$:

$$\frac{d^2\sigma}{d\Omega dk} = -\frac{|\vec{k}'|}{|\vec{k}|}\frac{G^2}{4\pi^2}\left(\frac{2\sqrt{2}}{g}\right)^2\int\frac{d^3r}{2\pi}\,\mathrm{Im}\left[L_{\nu\mu}\Pi^{\nu\mu}(q,\rho)\right]\Theta(q^0) \tag{1.62}$$

The W self-energy in a nuclear medium contains contributions from many possible loop diagrams such as those shown in Figure 1.12. Wavy lines represent the exchange boson, straight lines nucleons, double lines a Δ resonance, and dotted lines a strong interaction between nucleons. The contribution from each of these processes can be obtained through Cutkosky rules, in which the diagram is cut (dotted lines in the figure) and the cut particles are placed on-shell by taking their imaginary part. Each of these possibilities represents a potential final state. The process for subdiagram (1), in which two nucleons are placed on-shell, represents the simple quasi-elastic 1p1h interaction. Subdiagram (4) includes resonant pion production. But there are numerous other contributions, with subdiagrams (2), (3), (3') and (5) representing 2p2h processes, and (6) a 3p3h process.

Things get even more complicated when one considers that for each of these final states there may be several different vertices. Consider again subdiagram (2), representing a 2p2h process with a pair of nucleons coupled by the exchange of an intermediate meson (sometimes called Meson Exchange Current (MEC) events). The four-point vertex beneath the cut line may be expanded into one of seven possibilities, as shown in Figure 1.13.

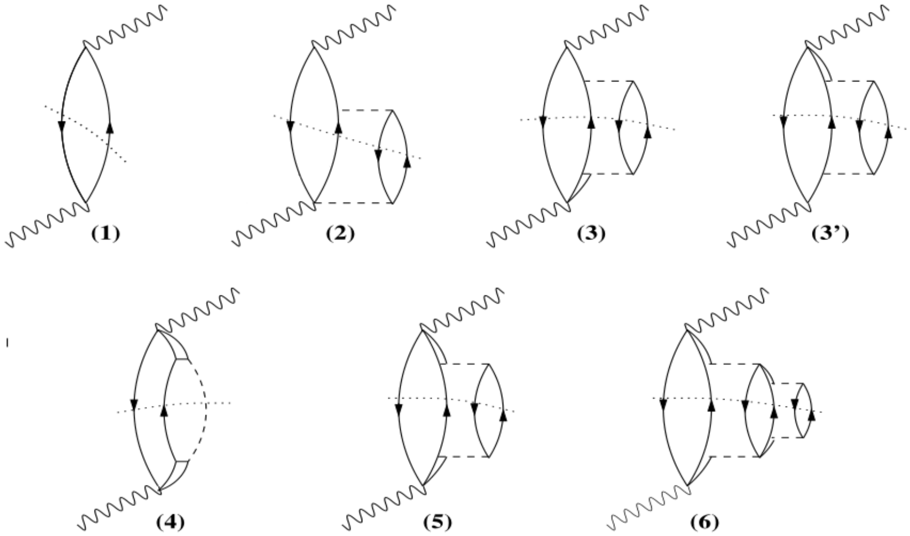

Figure 1.12: Some of the loop diagrams that contribute to the self-energy of the W boson in nuclear medium. Figure from [59]

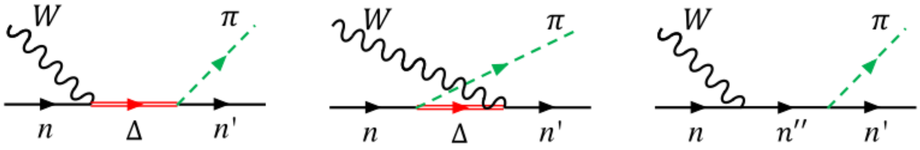

(a) Direct Δ pole
(b) Crossed Δ pole
(c) Direct nucleon pole

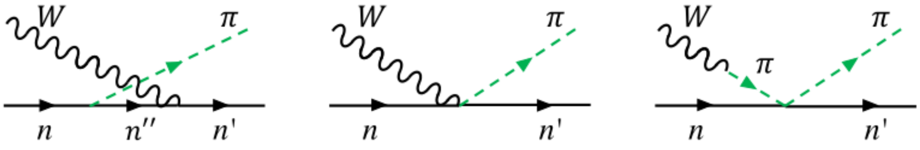

(d) Crossed nucleon pole
(e) Contact
(f) Pion pole

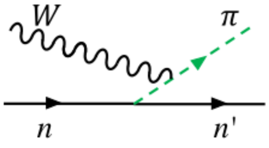

(g) Pion in flight

Figure 1.13: For the 2p2h process in Figure 1.12 diagram (2) there are 7 potential vertices that must be summed over. Figure from [50]

The primary impact of introducing 2p2h processes is a filling in of the ‘dip region’ between the QE and Δ resonance peaks in the $d\sigma/dQ^2$ cross-section (Figure 1.14).

As well as these microscopic approaches, empirical models were also developed to fill in the deficit observed in this region. These include the Dytman Empirical Model and the Transverse Enhancement Model (TEM).

The Transverse Enhancement Model[60][61] is based on extracting a parameterisation of the 2p2h enhancement to elastic electron-nucleus scattering cross-sections, and translating it to neutrino quasi-elastic scattering. The baseline electron scattering model is based on the scaling phenomenon. Scaling is a property of nuclear scattering cross-sections in which the free-nucleon cross-section scaled by the number of nucleons is found to match well with the total nuclear cross-section. Obtaining such a simple relation relies on finding a suitable ‘scaling parameter’ with which to parameterise the nuclear cross-section. An example of such a parameterisation can be found in ref [62]. Scaling is found to work well for the longitudinal part of the nuclear response function, but 2p2h processes result in an enhancement of the transverse component which must be accounted for independently. By fitting and extracting this enhancement, then applying it to predictions of neutrino quasi-elastic scattering, the Transverse Enhancement Model was an early success in reconciling MiniBooNE and global data without an inflated axial mass.

The Dytman Empirical Model[63] is a simpler model developed for modelling MEC contributions in the GENIE Monte Carlo event generator. It inserts a Gaussian contribution to the scattering cross-section between the QE and Δ resonance peaks, fitting the peak magnitude with electron scattering data and scaling linearly with nuclear size. The contribution grows initially with neutrino energy, but then falls off linearly to match higher-energy data.

1.4.6 Final State Interactions

To close off the discussion on neutrino-nuclear interactions we reconsider the plane wave approximation, namely that the ejected nucleon leaves the nucleus without further interaction. To assess the validity of this proposal, let us estimate the average number of interactions an ejected nucleon incurs on its way out of the nucleus by considering the schematic in Figure 1.15a. The ejected nucleon of momentum $\vec{p}$ traces out a cylinder from its initial starting point, with a radius based on the nucleon-nucleon cross-section $\sigma_{NN}(\vec{p})$. The number of scatterings is obtained by integrating the local nuclear density along this cylinder. Figure 1.15b shows the result for oxygen with $|\vec{p}| = 1\,\mathrm{GeV}$, starting from a perpendicular distance $b = 0.25\,\mathrm{fm}$, as a function of the initial axial location z. For $z < 0$ this approach gives at least one final state interaction

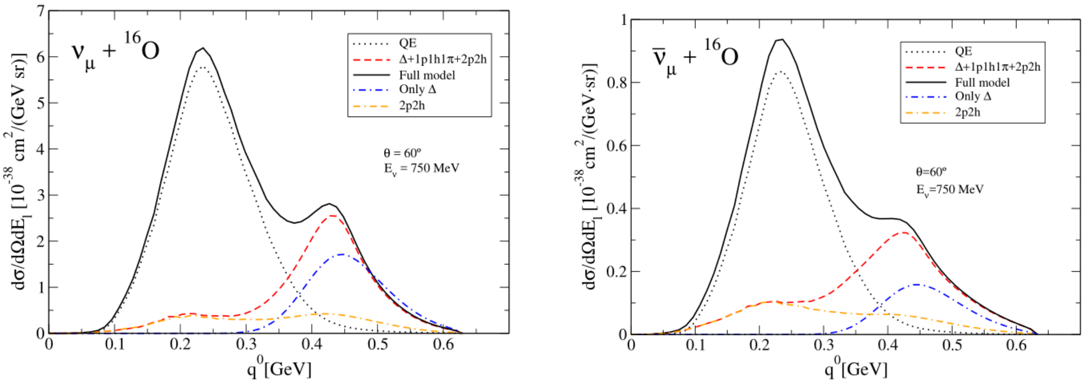

Figure 1.14: The inclusion of 2p2h interactions results in an enhancement of the differential CC cross-section in the ‘dip’ region, between the QE and Δ resonance peaks. Figure from [50]

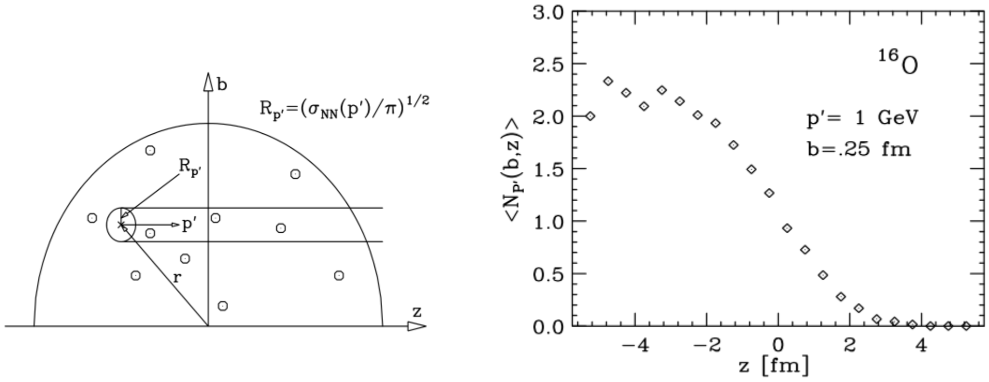

Figure 1.15: Left: Estimating the number of final state interactions for an ejected nucleon. Right: The resulting distribution evaluated under the shown parameters. Figures taken from [57]

on average, suggesting that the impulse approximation would not be justified. More rigorous treatments corroborate this conclusion[64]. These collisions can modify the energy of the outgoing nucleon or even knock out additional nucleons. There may be charge exchange interactions, where strong interactions change the charge of the emitted particle, or there may be absorption of pions produced in the initial interaction. All of these effects obfuscate the initial interaction and can result in an incorrect determination of the underlying event.

There are several approaches to modelling the effect of such Final State Interactions (FSI). One is to convolve the PWIA cross-section with a ‘folding function’ based on a complex potential[57][64]. Such a convolution can produce both a redistribution of events and an overall modification of the amplitude. Another approach is to introduce a particle spectral function, in a similar manner to the hole spectral function used to describe the momentum distribution of initial nucleons. The Heaviside step function that represents Pauli blocking is effectively one

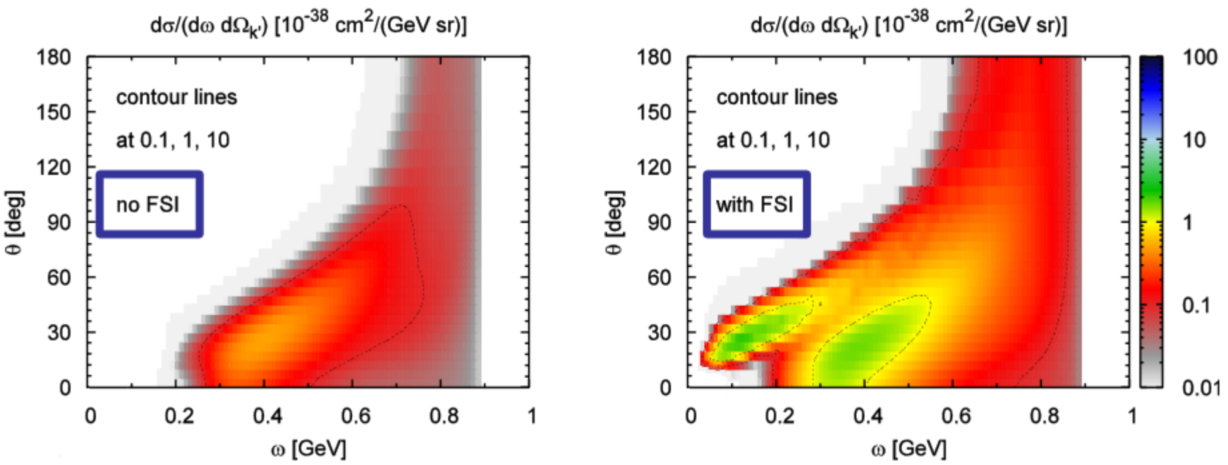

Figure 1.16: Charged current neutrino-induced neutron production from carbon, with and without FSI. Charge exchange interactions of the primary proton can result in a marked increase in neutron production. Figure from [67].

example of this. A more complete treatment may be performed by the same method used to account for long range correlations. Recall that the Random Phase Approximation accounted for nuclear-medium interactions of the exchange boson through the calculation of a dressed propagator. One can do the same calculation to derive a dressed outgoing nucleon propagator, including re-interactions of the nucleon as it propagates out of the nucleus, and a particle spectral function may then be obtained from the complex self-energy of this dressed propagator[49]. Further methods include the Glauber model[65][64][56], or computational transport models such as those in the neutrino event generators GiBUU or NuWro[66][39].

For inclusive cross-sections these effects are integrated out by summing over all interaction channels, but exclusive channels and differential cross-sections will be affected. For the differential quasi-elastic cross-section, $d\sigma_{QE}/dQ^2$, the result is only a slight broadening and a quenching of the peak amplitude. Other observables show greater changes - the energy spectrum of produced pions, for example, is shifted to lower energies with changes to both amplitude and shape. The generation of neutrons can also be significantly increased due to charge exchange interactions. Predictions for neutron knockout from carbon, calculated within the GiBUU framework, are shown in Figure 1.16 comparing results with and without FSI[67].

1.4.7 Neutrino-Nucleus Interactions Summary

The description of neutrino-nucleus interactions is a complex and challenging area, wanting in both theoretical and experimental inputs. While the broad range of available techniques is promising for our future understanding, the different approaches, approximations and notations can make it difficult to compare models on an even footing. Care must be taken when to ensure a consistent treatment of the initial state, nucleon correlations, multinucleon events and final state interactions, since the transition amplitudes for all processes producing a given final state will in general interfere. Without due care it is easy to incur double counting, over- or under-estimated systematic errors, or unexpected correlations. For more details on this, on the models discussed here, on others not discussed, and a look at how the available models compare, the interested reader is referred to one of the many reviews available: e.g. [68][36][69][37][70].

1.5 The Importance of Neutrino Cross-Sections

The knowledge of interaction cross-sections underpins all particle physics experiments, as it defines the rate with which a given interaction is expected to occur. Whether it's measuring oscillation parameters, determining the mass hierarchy, searching for CP violation or signs of sterile neutrinos, nearly all the major questions are answered by counting the number of interactions in some way or another. Interaction cross-sections are an essential pre-requisite to calculating the expected number of events, and therefore for getting answers to the questions at hand.

Measuring neutrino cross-sections is a very challenging task. All practical sources produce neutrinos via many-body decays of short-lived hadronic particles, which results in two major complications. Firstly, many-body decays produce a continuum of neutrino energies, and since the resulting neutrinos can't be manipulated with electric or magnetic fields, it isn't possible to refine this energy spectrum after generation. Determining the cross-section, which is a function of the neutrino energy, therefore relies on extracting the incident neutrino energy from the observed products. The previous sections hopefully illustrated that this is not straightforward to do.

The second challenge arises from the indirect nature of neutrino production. To achieve the high intensities needed for appreciable event rates huge numbers of neutrinos must be generated. This is done by bombarding protons onto a target, producing short-lived secondary hadrons that decay to produce the desired neutrinos. But the process of hadronic generation is itself a complex topic, requiring knowledge of hadronic production cross-sections and decay branching ratios. Uncertainties in these primary processes directly translate into uncertainties in the resulting neutrino flux, and since only the product of flux and cross-section is ultimately measured, this in turn leads to equivalent uncertainty on the neutrino cross-sections.

Together the theoretical and practical difficulties associated with determining neutrino cross-sections make them one of the largest contributions to systematic error in accelerator-based neutrino experiments. In T2K's 2016 paper on the observation of ν_e appearance[71], cross-section uncertainties accounted for 7.5% of the total 8.8% uncertainty on the predicted number of ν_e events. In the NOνA experiment cross-sections are again the dominant systematic on the far detector ν_e prediction, and are the second-largest systematic in the ν_μ prediction[72]. These uncertainties on the expected event rate translate into 35%, 44% and 53% of the total error on the extracted measurements of $\sin^2(\theta_{23})$, Δm^2_{32} and δ_{CP} respectively[73].

Next generation experiments are aiming for total uncertainties considerably lower than these levels. The Hyper-Kamiokande design report[74] projected a 3.2–3.9% systematic error on the number of events arising from combined uncertainties in the flux and cross-section, under the assumption that uncertainties from the use of different nuclear materials in the near and far detectors can be brought to negligible levels. The DUNE collaboration aims to bring the effect of cross-sectional uncertainties down even further, to below 2%[75]. However, a comparison using the event generators GENIE and GiBUU found that the sensitivity to CP violation varied by $\sim 1\sigma$ between the two models, the sensitivity to the mass hierarchy varied by as much as 6σ, and the sensitivity to the octant of θ_{23} varied by up to 3σ. These studies clearly demonstrate the importance of improving cross-sectional uncertainties and interaction modelling in achieving the desired physics goals.

ANNIE will be performing cross-section measurements to provide crucial data needed to support these improvements. As will be summarized in section 1.8, the extent of high quality neutrino cross-section data is currently very limited. Older experiments typically used light nuclei resulting in low statistics, and often measured only integrated cross-sections. More recently differential and double differential cross-sections in final state kinematics have started to emerge, and have brought to light significant deficits in current interaction models. Still, only a handful of such measurements exist for a very limited set of nuclear targets; the measurements obtained by ANNIE will be a valuable contribution to this dataset.

1.6 Measurement of Neutrino Cross-Sections

The interaction cross-section σ is derived from the rate of interactions N_{int} observed from a probe flux Φ onto a detector with target density T:

$$\sigma = \frac{N_{int}}{T\,\Phi} \tag{1.63}$$

During analysis the number of interactions is usually binned in the final state variables. The number of measured interactions in bin i will be a combination of true signal events S_i together with some background of events mis-identified as the desired type, B_i. In addition, not all true signal events will be measured, so we must correct by the efficiency associated with detection and reconstruction, ϵ_i. Finally, since the reconstruction of final state kinematics is not perfect there will also be some migration of events from true bin j into reconstructed bin i, which can be accounted for by a unitary smearing matrix U_{ji}. Altogether these give:

$$N_i^{int} = \frac{\sum_j U_{ji}(S_i - B_i)}{\epsilon_i} \tag{1.64}$$

We wish to use this number to estimate the cross-section for producing events in the kinematic region corresponding to this bin. Our kinematic variables are θ, the angle between the incoming neutrino and outgoing lepton, and E, the energy of the outgoing lepton. Given the neutrino beam divergence is $< 0.04\,\text{rad}$ we can take the neutrino as being parallel to the beam axis in the estimation of θ. The cross-section in this kinematic bin is then written as the average differential cross-section for producing events in the bin, $({}^{d^2\sigma}/_{dE_l d\theta})_i$, multiplied by the bin width ΔE_l, $\Delta\theta$. Combining this with Equation 1.63 and Equation 1.64 and rearranging for the cross-section gives:

$$\left(\frac{d^2\sigma}{dE_l\,d\theta}\right)_i = \frac{1}{T\int\Phi(E_\nu)\,dE_\nu}\frac{\sum_j U_{ji}(S_i - B_i)}{\epsilon_i\,\Delta E_l\,\Delta\theta} \tag{1.65}$$

It is worth noting that the measured cross-section includes an integration over the entire neutrino flux spectrum. This is because the mapping from initial to final state variables is not one-to-one. Experiments can only measure such 'flux-integrated' cross-sections; to make comparisons with theory one must convolve the predicted cross-sections (in terms of the initial state variables) with the experimental flux:

$$\frac{d^2\sigma}{dE_l\,d\theta} = \frac{1}{\int\Phi(E_\nu)\,dE_\nu}\int\left[\frac{d^2\sigma}{d\omega\,d\theta}\right]\Phi(E_\nu)\,dE_\nu \tag{1.66}$$

This convolution is generally done as part of the Monte Carlo event generation.
When comparing distributions between experiments it is important to be aware of the different regions of sensitivity introduced by this flux weighting. Consider the plots shown in Figure 1.17,

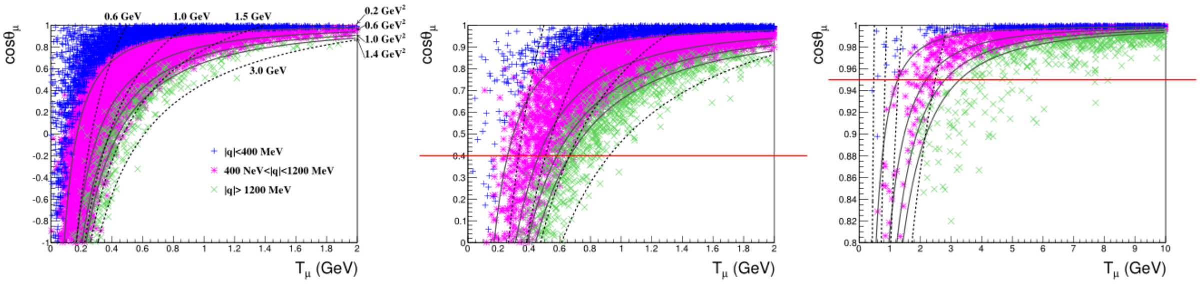

Figure 1.17: Distributions of true CCQE events on carbon, generated with the respective fluxes at the experiments MiniBooNE, T2K and MINERvA (left to right). Solid lines show solutions to free-nucleon scattering at fixed neutrino energy, dotted lines at fixed momentum transfer. Nuclear effects will smear the events abount these central lines. Horizontal red lines indicate detector acceptance. Figure from [76].

which show distributions of true CCQE events (Monte Carlo generated) in the measured E_l, θ plane. Each of these plots uses the neutrino flux from a different experiment. The dashed and solid lines represent solutions to the free nucleon scattering relations, Equation 1.56 and Equation 1.57, showing lines of constant neutrino energy and momentum transfer, respectively. By considering the density of events along these lines we can see that the different experiments will be sensitive to very different regions of the initial state variables.

1.6.1 Impact of Nuclear Effects on the Smearing Matrix

There are two primary impacts of nuclear dynamics on the reconstruction of neutrino events; the obfuscation of the underlying event type, and an incorrect inferral of the neutrino energy. In water Cherenkov detectors event reconstruction must rely on the fairly limited observation of charged particles above threshold; events can only meaningfully be described in terms of the number of these particles and their reconstructed kinematics. When making a selection of CCQE-like events the dominant contribution comes from true quasi-elastic interactions, but there are also significant contributions from 2p2h processes, pion production, and DIS events. Figure 1.18 shows the contributions to such a 'QE-like' cross-section as predicted by the GiBUU generator.

Equation 1.56 and Equation 1.57, on the other hand, rely on the assumption of a true quasi-elastic scattering. Their application to events that are not truly quasi-elastic tends to underestimate the true neutrino energy. Fits of theoretical predictions in terms of *true* neutrino energy to measured distributions in *reconstructed* neutrino energy then lead to incorrect fit parameters. This has been the subject of a number of publications[77][70][78][79], which warn that it can result in over-estimations of the axial mass, distorted cross-sections as a function of

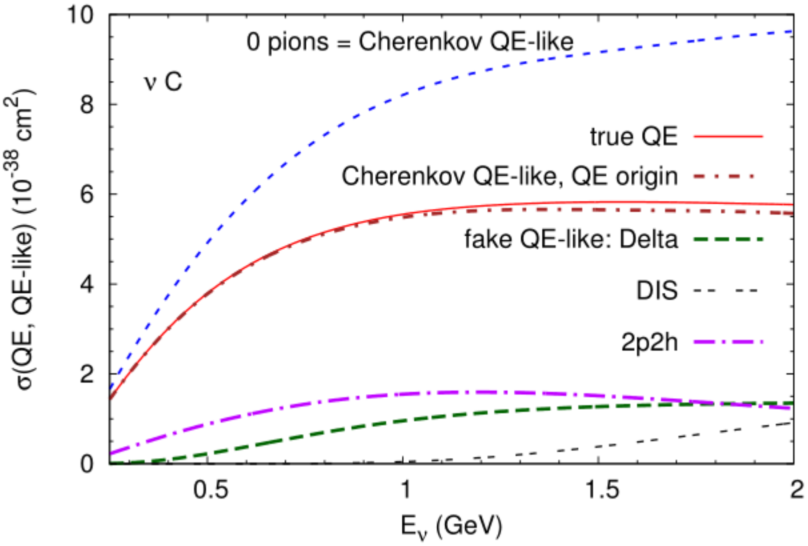

Figure 1.18: Different contributions to the total cross-section in QE-like events. Figure from [77]

energy, and a large systematic shift in derived oscillation parameters.
An example of this is illustrated in Figure 1.19, which shows the reconstructed neutrino energy for E_ν=1 GeV CCQE and QE-like events. The black line indicates all CCQE + QE-like events, the purple line shows true 2p2h interactions, and the red line shows pion production events in which the pion was re-absorbed within the nucleus; the remaining contribution is true CCQE. It is clear that the mis-labelled events skew the resulting distribution to lower values. This bias, which impacts the mapping from reconstructed to true energy, becomes a major contributor to the smearing matrix U_{ji} in Equation 1.64.

Traditional methods for handling this problem involved subtraction of the 'fake QE' components, or 'unfolding' processes that de-convolve the measured parameter space into that of the underlying interaction. But these methods are sensitive to the modelling of *all* contributions to the QE-like signal, and can introduce biases from mis-modelling of these contributions into the measured results. Nowadays it is preferred to publish results without unfolding[80][81], together with a set of detector response functions and correlated systematic error matrices. Although more cumbersome to compare against theory this avoids biases introduced in the unfolding procedure, allowing much more robust comparisons.

Given the broad consequences of event misclassification there is much interest in observables that provide discriminative power on the underlying interaction channel. One such handle is

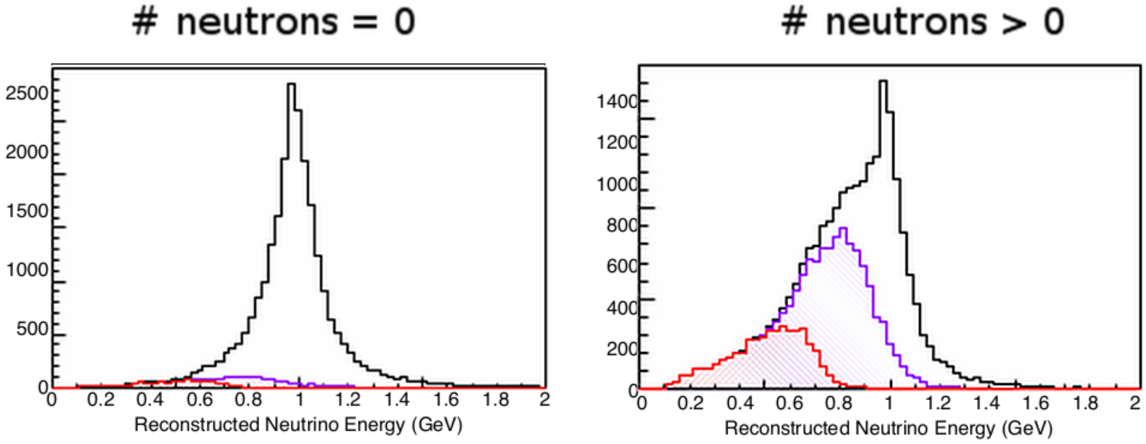

Figure 1.19: Energy reconstruction for 1GeV QE-like events, including (right) and excluding (left) those with neutrons in the final state. Figure from [82]

neutron tagging. Pion-nucleus scattering experiments suggest that nucleon knock-out from pion absorption may be dominated by neutron emission[83], suggesting neutron detection may be especially sensitive to inelastic collisions. A simulation of the potential impact on energy reconstruction is shown in Figure 1.19. On the left neutrino energy reconstructed using the CCQE assumption is shown for CCQE and CCQE-like events with no neutrons in the final state, while on the right the same plot is shown for events with at least one neutron in the final state. While not a perfect analogue of the true-CCQE vs CCQE-like divide, the presence of neutrons clearly provides a strong indication of inelasticity in the underlying event. As will be discussed in chapter 2, ANNIE will be applying a new technique to enable high-efficiency detection of final state neutrons, leveraging this discriminative power to obtain a minimally contaminated CCQE sample, and thereby producing cross-section measurements with minimal energy skew bias.

1.7 The Importance of Neutron Multiplicity

Neutron tagging in neutrino interactions has other physics applications beyond the selection of a purer CCQE sample; one example of particular interest is in searches for proton decay. Proton decay is a feature of many Grand Unified Theories[84] and its observation would be a historic discovery and a key piece of evidence in support of these theories. But despite decades of searches evidence for proton decay remains elusive, with current limits on the proton lifetime on the order of 10^{33} years[85][86]. Searching for such extremely rare events requires the strictest possible background suppression. Atmospheric neutrino interactions pose perhaps the greatest challenge as they cannot be suppressed, have a rate equivalent to a proton lifetime of 10^{31} yr and can closely mimic the signal presented by many proton decay channels. For experiments using water as the active medium the presence of neutrons in the final state can be an especially powerful discriminator. Proton decays in water are expected to produce a final state neutron in less than 6% of cases[87]. By contrast atmospheric neutrino interactions capable of mimicking proton decay signatures are expected to produce one or more neutrons in a majority of cases. A delayed neutron capture would therefore be a clear flag by which to tag neutrino backgrounds, substantially improving detection sensitivities. However, to assign discovery confidence to signals *without* a neutron requires a robust understanding both of detector efficiency and the expected neutron abundance from background events. The latter of these is currently based on simple interaction models that are known to be severely limited, so experimentally characterising the neutron abundance is a key pre-requisite to full utilization of this technique.

A similar argument can be applied to searches for supernova relic neutrinos (alternatively called the diffuse supernova neutrino background), which will be a key part of the up-coming Super-Kamiokande-Gd project[88]. Here neutron tagging will be used for identifying signal events representing $\overline{\nu}_e$ inverse beta decay reactions. Sub-Cherenkov threshold muons from atmospheric neutrinos are the dominant background in this search above 20 MeV. While these events are not expected to produce neutrons at tree level, an understanding of the fraction of background ν_μ reactions that do produce neutrons will be an important systematic uncertainty. Further applications of neutron tagging include statistical suppression of antineutrino backgrounds in neutrino events (at tree level $\overline{\nu}_\mu$-CCQE produces a neutron, while ν_μ-CCQE produces a proton), suppression of neutral current backgrounds in oscillation electron-neutrino appearance experiments (NCπ^0 and NCγ reactions are expected to produce more neutrons than ν_e-CC events), and providing additional information to assist in understanding a supernova burst should one be recorded. In all cases an understanding of the expected abundance is a key input to estimation of systematic uncertainties. For this reason the production of neutrons has been recognised as a critical area of study for future intensity frontier physics programmes[89][90].

1.8 Status of Neutrino Cross Section Measurements

Early measurements of neutrino cross-sections were performed with light nuclei bubble chambers, which provided low statistics but had minimal nuclear effects (both owing to the low A of the targets) and good detector resolution (typical of bubble chambers) enabling the extraction of high purity samples[91][92]. Antineutrino cross-sections were much less well measured due to the invisibility of the ejected neutron and lower still statistics. The Argonne National Laboratory (ANL) and Brookhaven National Laboratory (BNL) datasets are perhaps the most extensive and are often still used as reference points; for a recent re-examination of their results see [93] and references therein.

Some more recent past and present neutrino experiments are listed in Table 1.1, with published cross-sections given in Table 1.2[94]. To achieve greater statistics these experiments use a variety of heavier nuclei, with a few elements being of particular importance. Organic scintillator experiments are dominated by interactions on carbon, and correspond to the majority of published data. Oxygen is the main nuclear target for water Cherenkov detectors, but has surprisingly few cross-sections measurements. Cross-sections on iron have been measured through its use as a passive target in segmented tracking detectors. More recently argon has gained importance due to the rise of Liquid Argon Time Projection Chambers (LArTPCs).

The MiniBooNE experiment provides a good reference for the wealth of information that can be extracted from these high-statistics experiments. They have published cross-section analyses for CCQE-like, NCπ^0, CCπ^+, and NC elastic channels, among others[95]. They published the first double-differential cross-section measurements ($d^2\sigma/dT_\mu d\cos\theta_\mu$) in terms of final state topologies[58], along with background-inclusive data releases that allow model-independent comparisons with theory[96]. Two similar double-differential measurements have since been published by T2K[97][98]. Both experiments performed the measurements on carbon, with T2K's results coming from the ND280 detector.

Differential cross-sections ($d\sigma/dQ^2$) on iron have been published by NuTeV[99], and on liquid argon by ArgoNeuT[100, 101]. The MINERνA experiment is designed to measure cross-sections for a range of nuclear targets, with many results being presented as cross-section ratios for which many systematic errors will cancel. So far the experiment has published ratios of charged current inclusive cross-sections for iron, lead and carbon[102].

Oxygen has received surprisingly little attention from the experimental community, although it has been the subject of several theoretical papers using electron scattering data[103, 104, 105, 106]. A measurement of the axial mass for CCQE cross-sections on oxygen was performed by the K2K experiment[107], and two flux-averaged NCQE measurements have been presented by the T2K collaboration[108][109]. More recently differential measurements in muon momentum and scattering angle were also presented by the T2K collaboration[110], representing the first differential measurements on oxygen. Since the ND280 detector contains both water and carbon targets the published results included a comparison of the two cross-sections averaged over the shared flux. These suggested the two targets present very similar cross-sections, with the most notable difference being a larger cross-section on oxygen for high scattering angles.

Experiment	beam	$\langle E_\nu \rangle$, $\langle E_{\overline{\nu}} \rangle$ GeV	neutrino target(s)	run period
ArgoNeuT	$\nu, \overline{\nu}$	4.3, 3.6	Ar	2009 – 2010
ICARUS	ν	20.0	Ar	2010 – 2012
K2K	ν	1.3	CH, H_2O	2003 – 2004
MicroBooNE	ν	0.8	Ar	2015 –
MINERvA	$\nu, \overline{\nu}$	3.5 (LE), 5.5 (ME)	He, CH, H_2O, Fe, Pb	2009 –
MiniBooNE	$\nu, \overline{\nu}$	0.8, 0.7	CH_2	2002 – 2012
MINOS	$\nu, \overline{\nu}$	3.5, 6.1	Fe	2004 –
NOMAD	$\nu, \overline{\nu}$	23.4, 19.7	C	1995 – 1998
NOvA	$\nu, \overline{\nu}$	2.0, 2.0	CH_2	2010 –
SciBooNE	$\nu, \overline{\nu}$	0.8, 0.7	CH	2007 – 2008
T2K	ν, $\overline{\nu}$	0.6, 0.6	CH, H_2O	2010 –

Table 1.1: Some prominent past and present neutrino cross-section experiments

Experiment	inclusive	0π	$\pi^\pm$	π^0	ν_e
ArgoNeuT	CC	2p	CC	–	–
K2K	–	CC	CC	CC, NC	–
MINERνA	CC	CC, 1p	CC	CC	–
MiniBooNE	–	CC, M_A, NC	CC	CC, NC	–
MINOS	CC	M_A	–	–	–
NOMAD	CC	CC	–	NC	–
SciBooNE	CC	–	CC	NC	–
T2K	CC	CC, NC	–	–	CC

Table 1.2: Cross-sections measured by the experiments indicated in Table 1.1. All results except the last column are for $\nu_\mu/\bar{\nu}_\mu$ scattering measurements. For the corresponding publications see [94] and references therein

1.9 Status of Neutrino Multiplicity Measurements

Only two measurements of nucleon multiplicity have been performed to date.
Proton multiplicity measurements were performed by ArgoNeuT[111, 112], with the results presenting a significant discrepancy between the measured distributions and those predicted by Monte Carlo generators. This further serves to highlight the limitations of generator models, and a demonstration that multiplicity measurements provide another handle for quantifying nuclear effects. While the low energy resolution of LArTPCs allows the detection of protons down to a few 10s of MeV, the detection of neutrons is still more challenging. ArgoNeuT showed some sensitivity to neutron production through the detection of inelastic scattering of final state neutrons[113], but these measurements were not able to distinguish between neutrons and de-excitation photons. Even were such distinction possible these measurements still place a low energy threshold on the detected neutrons, which may result in an underestimation of low energy boil-off neutrons. Gadolinium doped water Cherenkov detectors offer a unique insight into nuclear effects in this respect.

Neutron multiplicity has been measured previously by the Super-Kamiokande collaboration, but the results have not been explicitly published[114]. The results used a dedicated low-energy trigger to search for neutron captures on hydrogen, whose 2.2 MeV decay gamma lies just on the threshold of visibility for the SK detector. An efficiency of just 20% was achieved, and the results are presented in terms of visible energy, combining ν_e, ν_μ, $\overline{\nu}_e$ and $\overline{\nu}_\mu$ samples. ANNIE should be able to considerably improve on these results by providing high statistics results exclusively for ν_μ, in terms of final state topology and with a good estimate of true neutrino energy.

2

The ANNIE Experiment

The Accelerator Neutrino Neutron Interaction Experiment (ANNIE) aims to measure interaction cross-sections and the multiplicity of neutrons generated from neutrino-nucleus interactions, as a function of final state kinematics. To this end the experimental detector consists of three main components:

1. A front scintillating counter to veto the products of interactions upstream of the detector
2. A water Cherenkov detector, doped with gadolinium sulphate and instrumented with Photo Multiplier Tubes (PMTs) and Large Area Picosecond Photo Detectors (LAPPDs), which acts as the neutrino target and provides information for event reconstruction
3. A downstream Muon Range Detector (MRD) consisting of alternating layers of steel and scintillator, to detect, stop, and measure the energy of final state leptons

A schematic of the ANNIE detector is shown in Figure 2.1. The stages of a typical event are shown in Figure 2.2, which proceeds as follows:

- Neutrino interactions occurring in the rock between the neutrino beam source and the water tank are tagged by the detection of charged products passing through the front scintillator wall. Such events are vetoed as they do not represent interactions in the water tank.
- The water tank provides the target for neutrino interactions of interest. Events originating within the tank are detected and reconstructed based on the pattern of Cherenkov light emitted by energetic charged products.

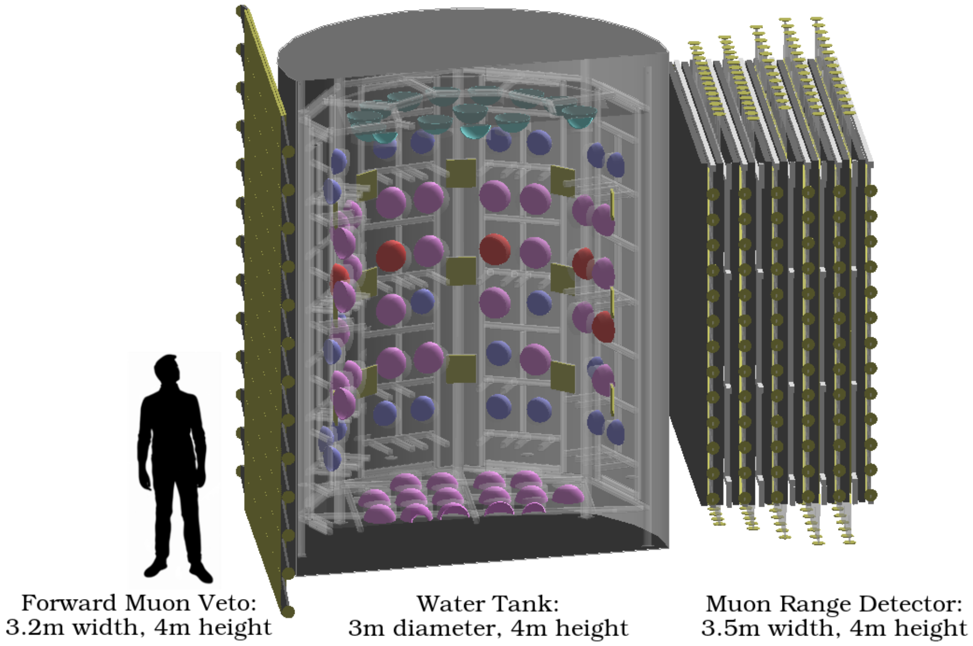

Figure 2.1: The ANNIE detector and its 3 main components. Beam enters from the left. At far left is the set of scintillator paddles that make up the front veto. The water tank in the center has been cut away to show the PMTs (hemispheres, colours represent PMT types) and LAPPDs (yellow squares). Right of the tank is the MRD. Image exported from the WCSim Monte Carlo framework.

- Primary muons leaving the water tank with sufficiently forward angle will enter and be slowed by the MRD. Muons that stop may have their energy determined from their penetration depth. The energy and angle of the muon is used to reconstruct the neutrino energy and momentum transfer.

- Any final state neutrons will thermalise through scattering with protons in the water, eventually capturing on a gadolinium (or possibly hydrogen) nucleus within a few 10s of centimetres from the interaction vertex.

- The excited nuclei will then decay through the emission of gamma rays, producing delayed flashes of light that are counted by the photodetectors. This is the mechanism by which neutrons are counted.

Table 2.1 gives an summary of key specifications of the detector; each component will be described in more detail in later sections.
A timeline of the experimental progress so far is given in Table 2.2. As will be discussed in subsection 5.4.2 an initial ‘Phase I’ background measurement was performed between 2015 and 2017 using a partially instrumented detector. This initial phase confirmed that the amount of ‘sky-shine’, a source of background neutrons that is difficult to estimate from Monte Carlo

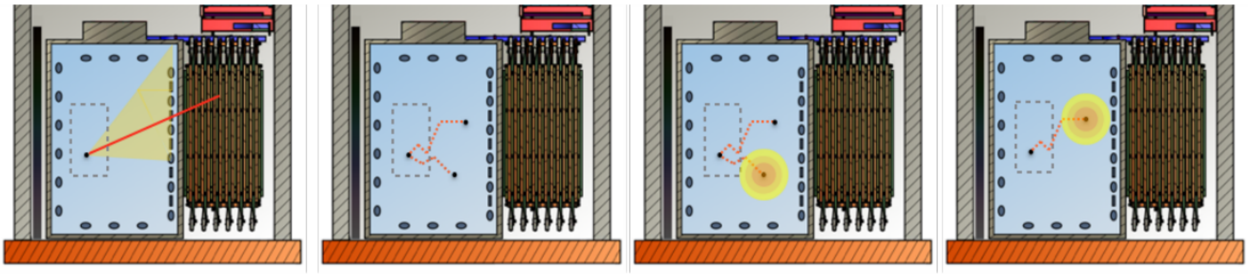

Figure 2.2: Stages of a typical event in the ANNIE detector. 1) A muon neutrino interacts in the water tank, producing a muon that emits a cone of cherenkov light before exiting the tank and stopping in the MRD. 2) Neutrons thermalise in the water. 3) & 4) Neutrons capture on gadolinium nuclei, which undergo subsequent decay generating delayed flashes of cherenkov light. Figure from [82]

studies, was sufficiently low to support the planned physics programme. Results from this study demonstrated that sky-shine levels are indeed well below the required levels[115]. Based on these results the full physics programme was approved in fall 2017 and the detector was upgraded to its full design specification over the following 18 months. More recently ANNIE has undergone commissioning and begun physics data taking. This is expected to continue for a 2-year period, after which there may be considerations for future upgrades to expand the physics programme. The work in this thesis was performed over a period spanning Phase I construction through to Phase II commissioning.

Description	**Value**
Total Water Volume	26 tonne
Fiducial Water Volume	2.5 tonne
Peak of Neutrino Energy Spectrum	$\sim 700\,\mathrm{MeV}$
Neutrino Fluence	$4.4 \times 10^{12}\,\mathrm{cm}^{-2}\,\mathrm{yr}^{-1}$, integrated over all energies
Rate of Neutrino Interactions	$\sim 38,000\,\mathrm{yr}^{-1}$ within the fiducial volume
Stopping Power of MRD	$800 - 1200\,\mathrm{MeV}$, dependent on penetration angle
Water Tank Photosensor Coverage	$\sim 8\,\%$
Event Vertex Resolution	68 % reconstructed within 12 cm
Primary Muon Energy Resolution	$\sim 6\,\%$
Neutron Capture Detection Efficiency	$\sim 30 - 60\,\%$, position dependent
Background Neutron Rate	0.06 ± 0.02 per beam spill, within the active volume

Table 2.1: Key specifications of the detector. Resolutions and efficiencies are preliminary.

Operation	**Start Date**	**End Date**
Phase I Approval	-	Spring 2015
Phase I Planning	Spring 2015	Winter 2015
Phase I Construction	Winter 2015	Spring 2016
Phase I Commissioning	Spring 2016	Summer 2016
Phase I Data Taking	Summer 2016	Fall 2017
Phase I Publication, Phase II Approval	Fall 2017	Spring 2018
Phase II Construction	Spring 2018	Summer 2019
Phase II Commissioning	Summer 2019	Spring 2020
Phase II Data Taking	Summer 2020	Fall 2022

Table 2.2: Timeline of the experiment.

2.1 Forward Muon Veto (FMV)

The ANNIE forward veto consists of 26 scintillator paddles obtained from the decommissioned CDF detector[116]. Each paddle consists of a 322 cm x 31 cm x 2.5 cm rectangular active area, coupled to a 36 cm long tapered glass light guide at one end. The light guide narrows to 5 cm where it is connected to a 2-inch (5 cm) diameter EMI-9815 PMT. The paddles are arranged in two horizontal layers, shifted vertically by 0.5 cm with respect to each other to ensure there are no uncovered gaps. In total this produces a 4 m high by 3.2 m wide wall that completely covers the beam-facing detector acceptance.

Measurements taken using a 3-fold coincidence found a detection efficiency of $\sim$90% for cosmic muons penetrating paddles near to the PMT, but this fell to 30-60% at the far end of the paddle, suggesting a reduction in transparency due to scintillator ageing. In order to maximize the detection efficiency the two layers are therefore combined in an 'OR' configuration, rather than the more usual 'AND' used to reduce the rate of false positives due to PMT dark noise. The 'OR' configuration significantly improves the efficiency of detection, with the result describing a parabola extending from $\sim$95% at the edges to $\sim$90% at its lowest point in the centre of the wall. The measured beam-off discriminated pulse rates of $\sim$200 Hz are in good agreement with the expected cosmic muon rate, suggesting the pulse discrimination threshold is sufficiently high to eliminate PMT dark noise. Given the short beam windows cosmic veto hits are expected to result in false positives in $\sim$ 20% of beam spills in this 'OR' combination, which is considered an acceptable loss. A precise measurement of this rate will be established once final voltages are tuned, as it will be important for any cross-section measurements. A photograph of the forward veto is shown in Figure 2.3.

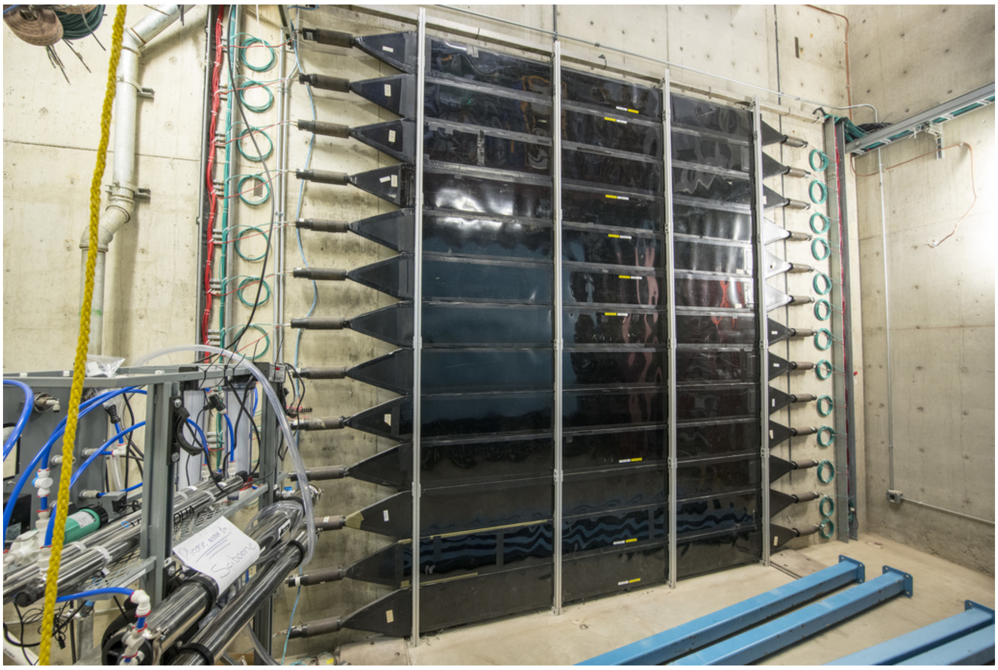

Figure 2.3: The ANNIE forward muon veto

2.2 Water Cherenkov Detector

2.2.1 Cherenkov Radiation

When light propagates through a transparent medium its speed is slowed to $c' = {}^{c}/_{n}$, with c the speed of light in vacuum and n the refractive index of the medium. As a result it becomes possible for a sufficiently energetic particle to move faster than light within that medium. As first observed by Cherenkov and Vavilov[117], and later explained by Tamm and Frank[118], such superluminal charged particles will produce a continuous spectrum of electromagnetic radiation. The particle itself is not the source of the emission; rather, as the particle travels it briefly polarises the surrounding medium as visualized in Figure 2.4. As these excitations relax the electric field disturbances interfere, and when $v > c'$ a wavefront forms as depicted in Figure 2.5. The wavefront propagates at a characteristic angle defined by $\cos\theta_c = {}^{c}/_{vn}$. For $\beta = {}^{v}/_{c} = 1$ and $n = 1.33$ this gives an initial emission angle of $\theta_c = 42°$, defining a cone that collapses as the particle slows down. The energy thresholds in water from the criterion $v > c'$ are given for some common particles in Table 2.3.
The number of photons emitted per unit length by a particle of charge Ze as a function of wavelength λ is given by Equation 2.1[119].

$$\frac{d^2N}{dxd\lambda} = \frac{2\pi\alpha Z^2}{\lambda^2}\left(1 - \frac{1}{\beta^2 n^2(\lambda)}\right) = \frac{2\pi\alpha Z^2}{\lambda^2}\sin^2\theta_c \tag{2.1}$$

Here α is the fine structure constant. The dependence on λ^{-2} results in a greater intensity at shorter wavelengths, although this is generally curbed beyond the UV region as the refractive index tends to 1. This wavelength dependence is the origin of the characteristic “blue glow” often seen in photographs of nuclear reactor cores. The dependence on β^{-2} demonstrates that the rate of photon emission is only weakly dependent on the particle energy, so the number of photons emitted per unit track length is roughly constant. (The lower limit $\beta = 1/n$ and upper limit $\beta = 1$ means the rate of photon emission only changes by a factor of 2 from threshold to infinite energy). For an ultra-relativistic particle with unit charge traversing through water this amounts to ∼300 photons per cm in the range of 300–600 nm. This represents less than 0.1% of the particle’s total energy loss (the majority goes to ionization and atomic excitation), or around 100 times fewer photons per MeV than in typical scintillators. High transparency and good photon detection efficiency are therefore essential features of a water Cherenkov detector.

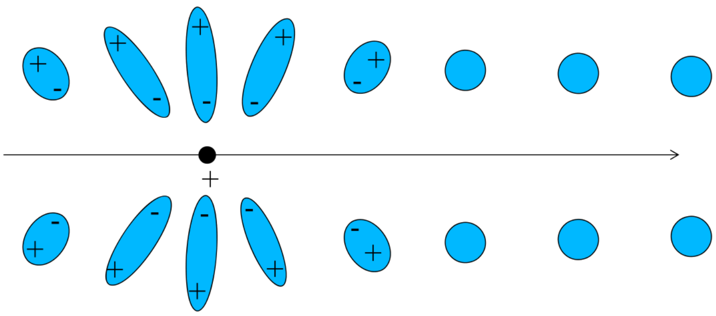

Figure 2.4: Visualization of the polarization of the surrounding medium induced by the passage of a charged particle. Figure from [120].

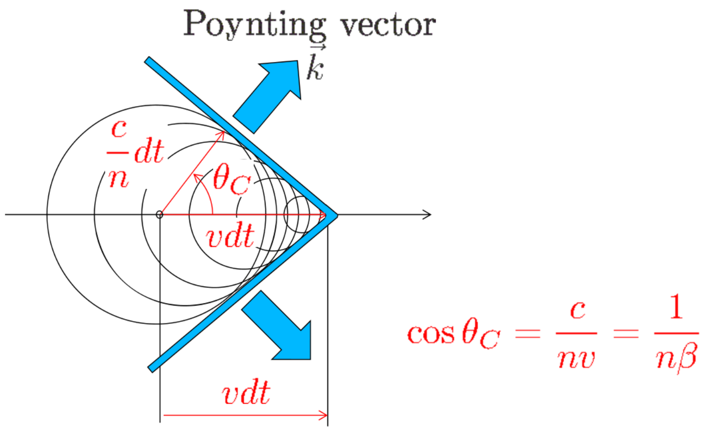

Figure 2.5: When the particle velocity v exceeds the local speed of light c, constructive interference of field disturbances results in the formation of a free wave, indicating the presence of radiation emission. Figure from [120].

2.2.2 The ANNIE Water Tank

Water Cherenkov detectors have a long history of use in neutrino detection, from past experiments such as IMB[121], SNO[122] and Kamiokande[123]; to present ones including Super-Kamiokande[124], IceCube[125] and ANTARES[126]; and even to future experiments such as Hyper-Kamiokande[127] and KM3Net[128]. Standing just 4 m tall and with a 3 m diameter, the ANNIE detector is perhaps one of the smallest water Cherenkov neutrino detectors ever built. This poses both advantages and disadvantages. Water transparency requirements are relaxed since the maximum travel distance is much smaller, with proportionally less stringent requirements on the water purification system. This is particularly useful given the additional challenges associated with gadolinium loading. On the other hand vertex reconstruction is much more challenging since the angular granularity of photosensor coverage is worse, edge effects are

Particle	Threshold Total Energy [MeV]
$e^{\pm}$	0.768
$\mu^{\pm}$	158.7
$\pi^{\pm}$	209.7
p^{+}	1400

Table 2.3: Total relativistic energy thresholds for Cherenkov radiation from common particles in water.

more pronounced, and fiducialization cuts must be as conservative as possible.

The water container is made from 7-Gauge (4.5 mm) steel to provide the detector structure. The steel is not stainless but is lined with several layers of plastic so is not in direct contact with water. Two initial layers of 6 mil (0.006") polyethylene and a 30 mil bottom sheet provide protection from rough areas in the steel. The liner that provides primary containment is made from 40 mil PVC #328 NSF-61, manufactured by Kentain[129]. Particular care was taken to ensure no UV stabilizers were present in the liner formulation, since these can leach into the water and absorb the Cherenkov light - the chosen material uses a diisononyl phthalate plasticizer and zinc-calcium thermal stabilizer, with no additional UV stabilizer.

2.2.3 Photo Multiplier Tubes (PMTs)

Within the water tank 132 PMTs are mounted on a stainless steel support structure, amounting to 8% photocoverage for events at the centre of the tank. The majority of these PMTs are repurposed from previous experiments, resulting in a variety of different types being used, but all operate on the same basic principles. The main body of the PMT is a borosilicate glass bulb with an evacuated chamber at the front and a series of staggered dynodes at the back (Figure 2.6). The front face of the bulb is activated by a thin bialkali coating with a low work function, such that incident photons of visible energies are able to free electrons with moderate efficiency. This front photocathode is held at a potential of ∼-2 kV, which accelerates the freed electrons towards the dynode chain at the back. On impacting the first dynode each accelerated electron has enough energy to free 3–5 more electrons. Each dynode is held at a progressively lower voltage and the electron cascade is accelerated from stage to stage, producing an avalanche of some 10^7 electrons by the time the final dynode is reached. This leaves the PMT as a current pulse of a few nanoseconds, with this timing resolution dominated by the 'transit time spread' (TTS) due to the variation in electron path lengths. As will be shown in section 6.3 this timing resolution is a key property for Cherenkov detectors as it directly translates into the positional resolution of the reconstructed vertex.

The efficiency of detection is a product of the ‘quantum efficiency’ (QE) with which electrons are ejected from the photocathode on receipt of a photon, and the ‘collection efficiency’ with which this electron is properly accelerated through the avalanche path. The wavelength dependence of the quantum efficiency for a Hamamatsu R5912 PMT is shown in Figure 2.7, along with the shape of the Cherenkov emission spectrum and their product. The greatest contribution to the detection efficiency is around 330 nm, so transparency in this region is critical.

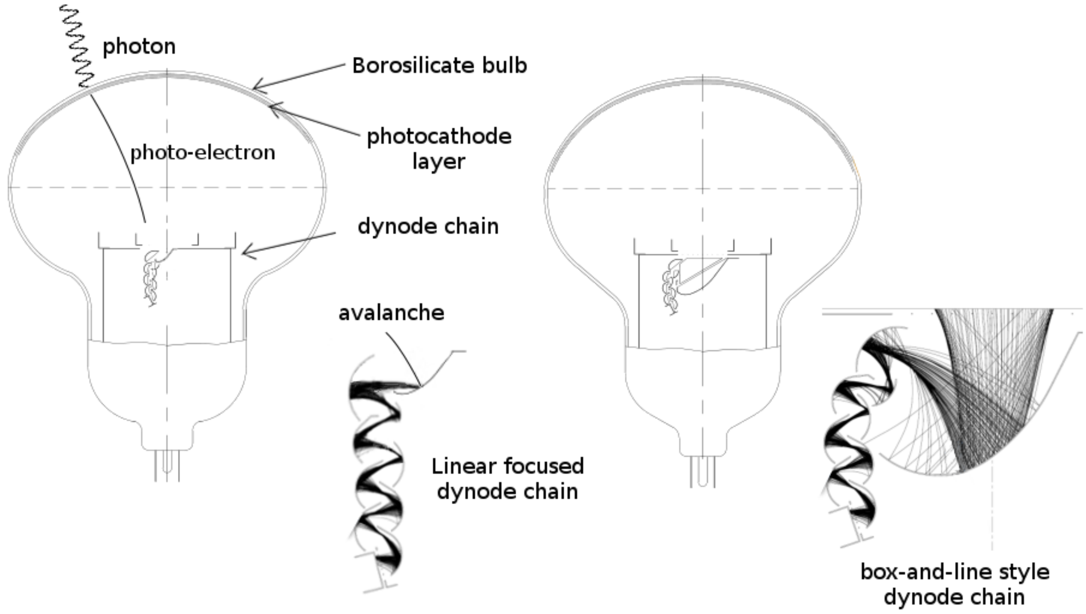

Figure 2.6: Schematic of a PMT. Left and right show two types of dynode chain in ANNIE PMTs. Figure adapted from [130].

Returning to the ANNIE detector, there are 4 specific models of PMT in the water tank. Twenty Hamamatsu R7081 PMTs, previously used in the LUX experiment[132], are arranged in a rectangular grid on the bottom of the tank. These PMTs are 10” in diameter, with a 10 stage box-and-line dynode, a nominal quantum efficiency of 25% at 390 nm, and a datasheet gain of $1.0\text{x}10^7$. The top cap is covered by twenty prototype ETEL D784UKFLB PMTs, initially used in a characterisation study for the LBNE project[133]. The tubes are 11” in diameter, with a twelve-stage linear focused dynode, a quantum efficiency of 32% at 390 nm, and a nominal gain of $1 \text{ x } 10^7$. Due to the presence of a circular hatch in the top cap these PMTs are arranged in an outer ring of 16, with the remaining 4 PMTs mounted in a central square on the underside of the hatch.

The main body of the support structure is octagonal in shape, with 6 rings of PMTs, 2 on each face, populated with 3 different PMT types. The most numerous are Hamamatsu R7081 PMTs, forty eight of which were obtained from the decommissioned WATCHBOY detector[134]. These PMTs are the same type as those used in the LUX experiment, with characteristics described above. A further forty Hamamatsu R5912-100 PMTs were purchased new by the ANNIE collaboration. These PMTs are 8" in diameter, with a 10 stage box-and-line dynode, a quantum efficiency of 35% at 390 nm, and a nominal gain of 1 x 10^7. Finally, four additional R7081-100 PMTs were obtained on loan from the WATCHMAN collaboration[135]. These are identical to the R7081 models used in the LUX and WATCHBOY experiments, but with an improved photocathode formulation giving a higher quantum efficiency of 35%. This gives a total of ninety-two barrel PMTs, leaving four remaining slots in the bottom ring of the support structure that are not filled. To minimize reflections that interfere with vertex reconstruction, and to optically isolate the active volume from the buffer region around it, the inner structure is surrounded by a black 4 mil (0.004") low density polyethylene sheet, again with no UV stability additives. A photograph of the inner structure, populated with the PMTs, is shown in Figure 2.8.

Along with the PMTs ANNIE will use LAPPDs to detect Cherenkov light with finer resolution, improving reconstruction of the event. LAPPDs will be discussed in more detail in section 2.11. For the first physics stage of the experiment five LAPPDs will be deployed in the tank, with a plan to increase this number as more become available in the future.
The LAPPDs are entirely enclosed within a 17"x13"x2" waterproof housing, consisting of a PVC frame, a UV-transparent acrylic front window and a stainless steel back plate; an LAPPD in one of the housings is shown in Figure 2.9. In order to allow the installation of new LAPPDs without needing to drain the tank and remove the support structure, a set of eight 304 stainless steel Unistrut rails are incorporated into the support structure, one at each corner. The LAPPDs are then mounted onto a 1/4" PVC panel, featuring four polypropylene sliders fixed to its back. The panel is then directed, in-situ, onto the Unistrut rails through one of 8 'mail slot' hatches in the tank lid. Each mail slot is capped with an opaque PVC lid featuring a slot for cable feedthrough. The setup is shown in Figure 2.10.

Readout of the LAPPDs is incorporated into the DAQ hierarchy as shown in Figure 2.16, with the ANNIE Central Card taking the place of the ADC card in interfacing with the LAPPDs. The details of the LAPPD communications system will be discussed in section 2.11.

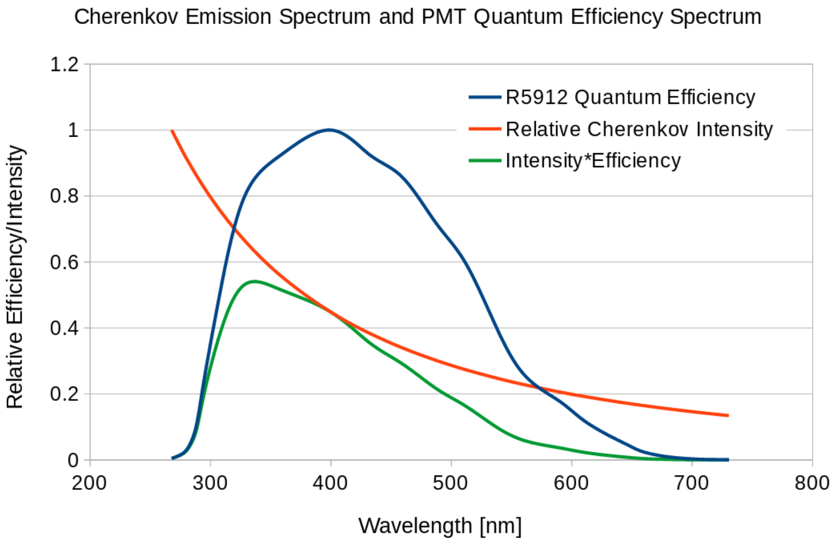

Figure 2.7: Comparison of the emission spectrum of Cherenkov radiation and quantum efficiency of a Hamamatsu R5912 PMT. The Cherenkov spectrum is normalized to 1 at the minimum plotted value of 267 nm, the PMT quantum efficiency is normalized to 1 at its peak efficiency at 400 nm. The product is also plotted, showing the relative contribution to the detected signal of each wavelength. QE data was extracted from the datasheet graph using WebPlotDigitizer[131].

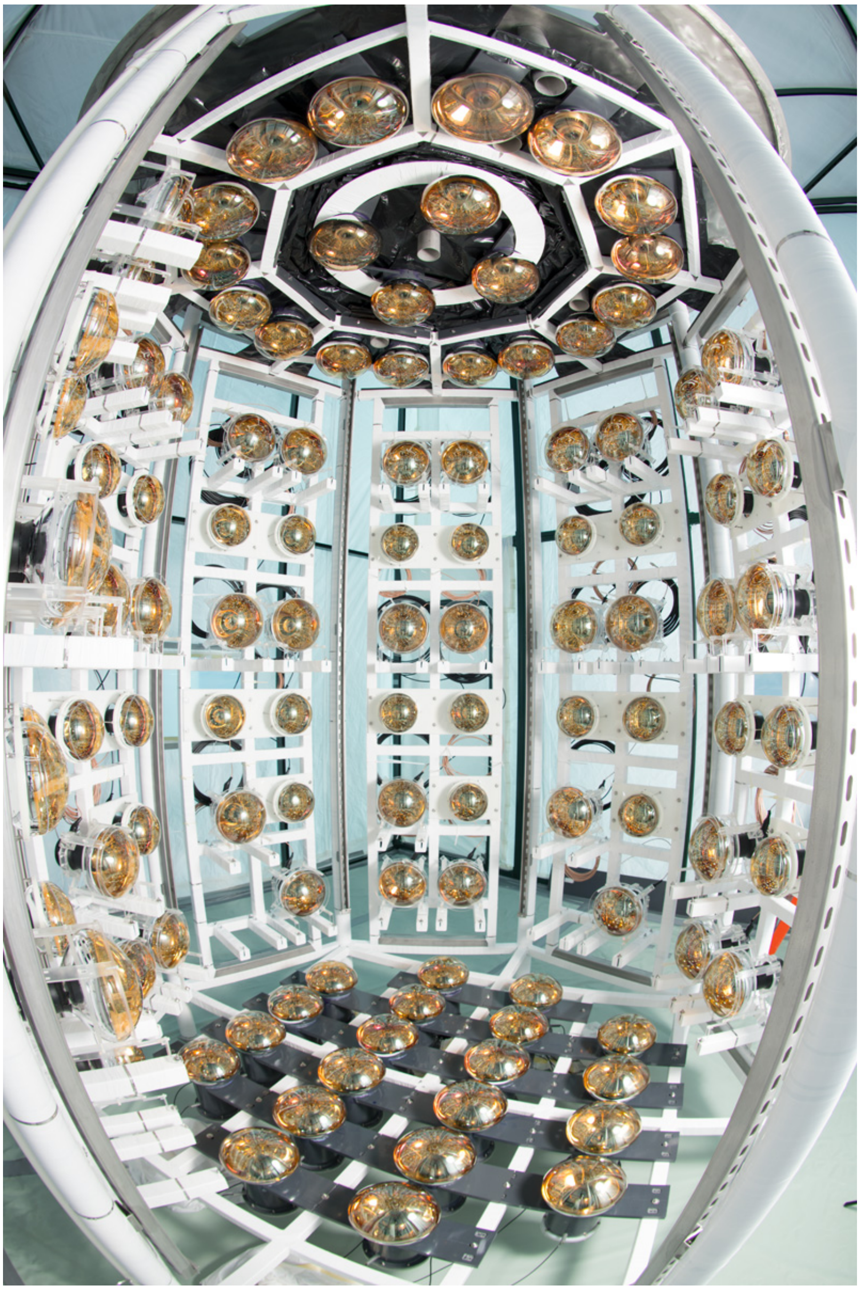

Figure 2.8: A photo of the inner structure, populated with PMTs, before installation into the water tank. At this stage the optical isolating blacksheet is installed in the top cap only. Note the photograph used a fish-eye lens.

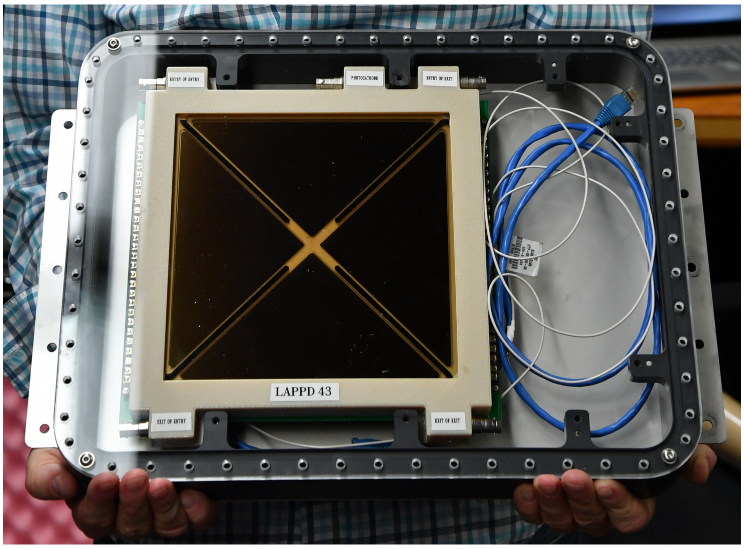

Figure 2.9: An LAPPD within its waterproof enclosure. The empty space to the right of the LAPPD will hold the electronics board, providing high voltage generation and LAPPD readout.

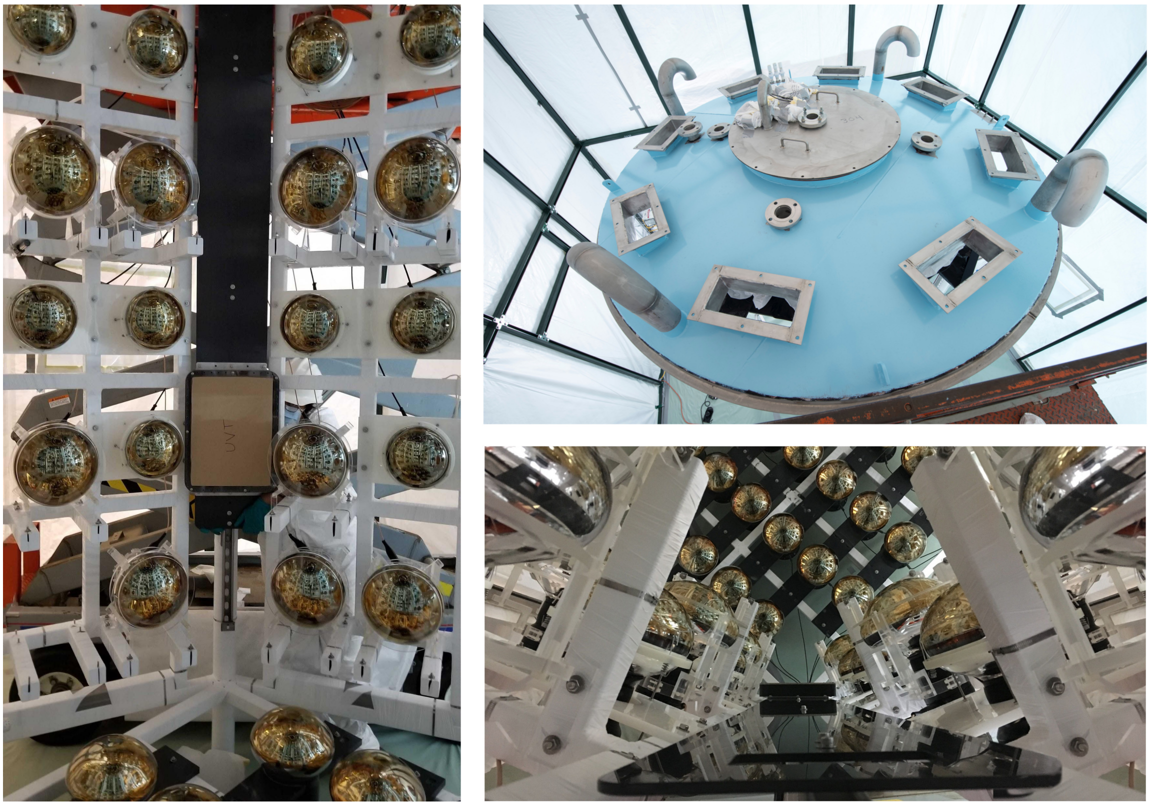

Figure 2.10: Left: A panel with one LAPPD housing being fit-tested on the unistrut slider. Polypropylene blocks on the back of the panel fit into the unistrut rail, the end of which may be seen below the LAPPD housing. Top Right: The 'mail slot' hatches on the tank lid. Bottom Right: A view looking down one of the 'mail slot' hatches.

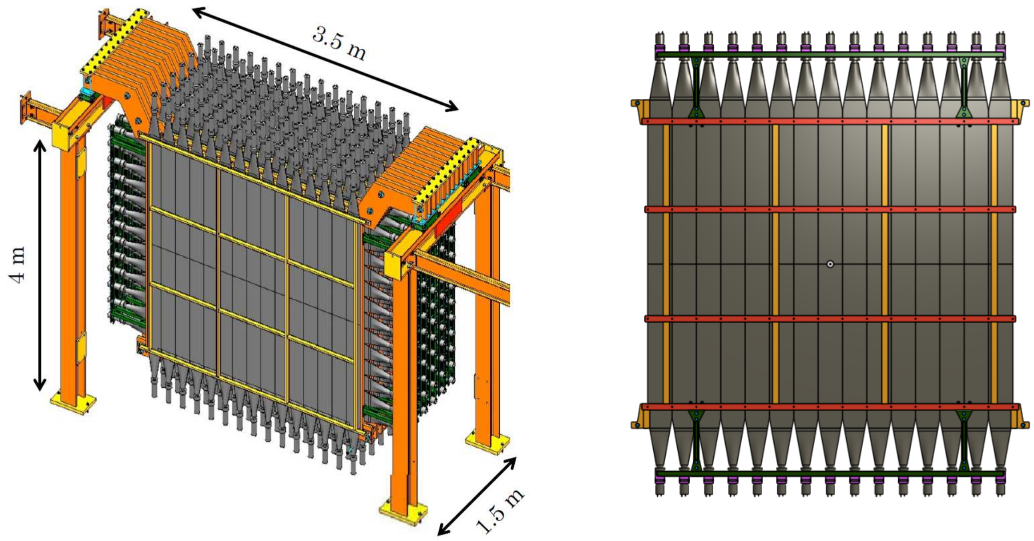

Figure 2.11: Left: An isometric view of the original SciBooNE MRD. Right: A view of one of the vertical scintillator layers. The orange & red grid overlaid shows the aluminium frame that supports the paddles.

2.3 Muon Range Detector

The Muon Range Detector (MRD) was inherited from the SciBooNE experiment[136] which ran from 2007–2008 in the hall that ANNIE now occupies. A schematic of the original MRD is shown in Figure 2.11, left. In order to maximize the space available for the ANNIE tank the front steel absorber and two scintillator layers were removed, leaving 11 steel absorbers and 11 layers of scintillator. Each steel absorber is 274 cm x 305 cm x 5 cm in size with an average density of 7.84 g/cm^3. This results in a stopping power of ∼800 MeV for muons at normal incidence, but since tracks at an angle can see a greater depth of steel the total MRD stopped sample actually spans energies up to ∼1.1 GeV.

The original scintillator layers alternate between 30 paddles in horizontal orientation and 26 paddles in vertical orientation, with the paddles in each layer configured as two sets placed end-to-end (see Figure 2.11right). Horizontal layers are instrumented with EMI 9954KB PMTs with a datasheet gain of 1x10^7, while the vertical layers use RCA 6342A PMTs that have a lower gain of 1x10^6. In order to provide sufficient signal for the CAMAC discriminators these PMTs are amplified by 10x using NIM amplifiers.

Upon SciBooNE's decommissioning 71 scintillator paddles were removed from the vertical layers. To restore the detector 64 replacement paddles were acquired from the past KTeV experiment[137]. These paddles are slightly narrower, resulting in a closer paddle spacing and

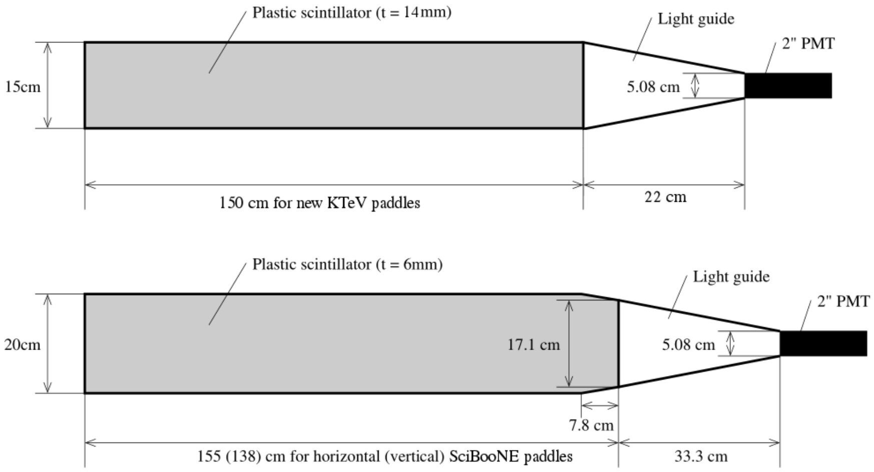

Figure 2.12: Dimensions of the MRD scintillator paddles.

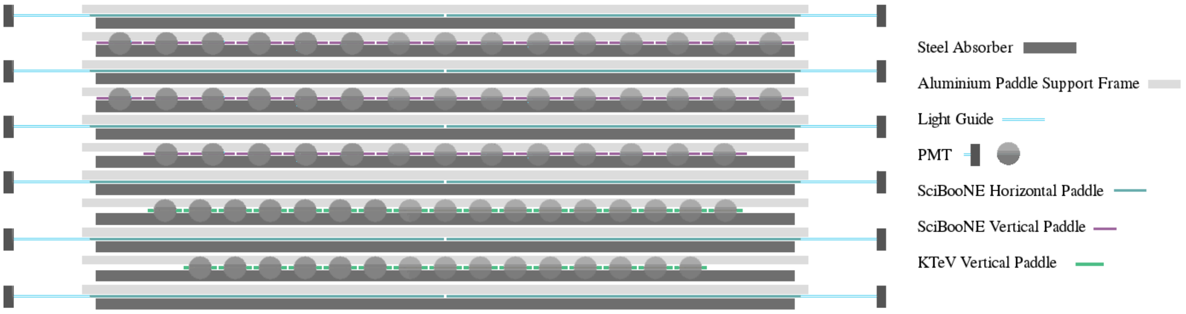

Figure 2.13: A top-down view of the MRD, showing the arrangement of scintillator paddles; beam enters from the bottom. The front two layers use narrower KTeV paddles, while the back three layers use the old SciBooNE paddles. All horizontal layers use the original SciBooNE paddles.

a slightly reduced span in the first few layers. The new paddles are instrumented with EMI 9954KB PMTs, the same as those for the horizontal layers, so do not need amplification. The dimensions of the paddles are shown in Figure 2.12, with the resulting configuration of vertical layers shown in Figure 2.13.

2.4 Outer Detector

Although not a part of the final design, an outer detector was considered to assist with identifying particles entering the inner volume from interactions outside of it. The proposal considered sixteen 2-inch PMTs mounted outside of the blacksheet-wrapped active volume, one positioned at the top and bottom of each face of the inner structure, pointing toward the center (Figure 2.14). The white liner of the water tank would help to maximize reflections and improve detection efficiency, while the blacksheet wrapping around the inner structure would prevent light from leaking into the active volume and impacting reconstruction. The efficiency of this setup was estimated with WCSim, using primary muons originating from neutrino interactions upstream of the detector ('dirt muons'). The results are shown in Figure 2.15, where successful detection was defined as having at least two OD PMTs with a hit. The detection efficiency depends strongly on muon energy, rising from 5% at 400 MeV to nearly 30% at 1200 MeV. Although the dirt muon energy distribution peaks at ∼400 MeV, requiring at least 2 penetrated layers in the MRD would pull the distribution of selected muon energies up, resulting in a peak at around 800 MeV. Muons meeting this criteria, and thus presenting a background for fully reconstructed charged-current events, could therefore be expected to see an OD efficiency of at least 15%. [1] These results were considered by the collaboration, but ultimately owing to the low efficiency and time constraints, the outer detector PMTs were not included in the final installation.

---

[1] Low energy dirt muons could potentially contaminate the tank reconstruction if they coincided with a true neutrino event, but the rate of such accidental coincidences would be extremely small.

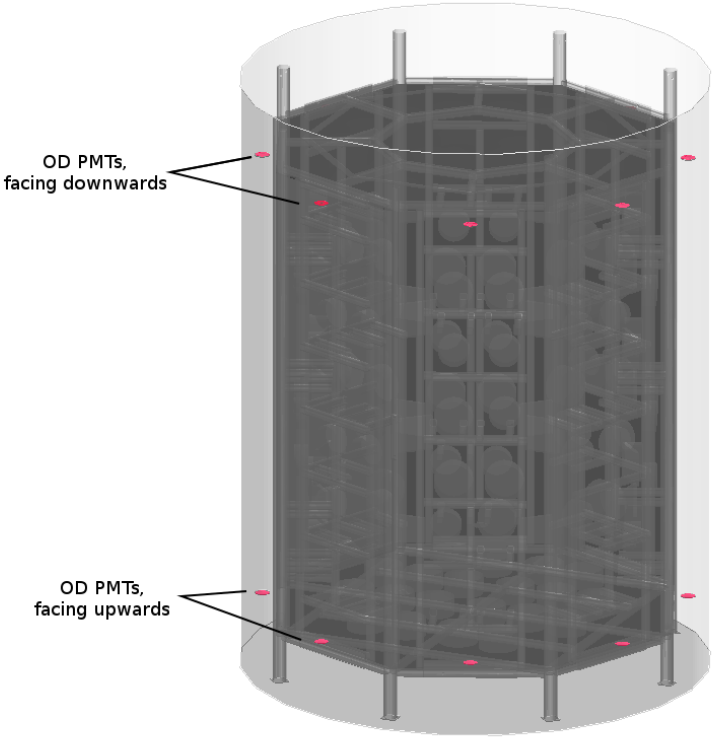

Figure 2.14: A schematic of the water tank highlighting the locations of the proposed outer detector PMTs. The PMT faces are shown in red; the body of the PMT (not shown) would be mounted to the inner structure.

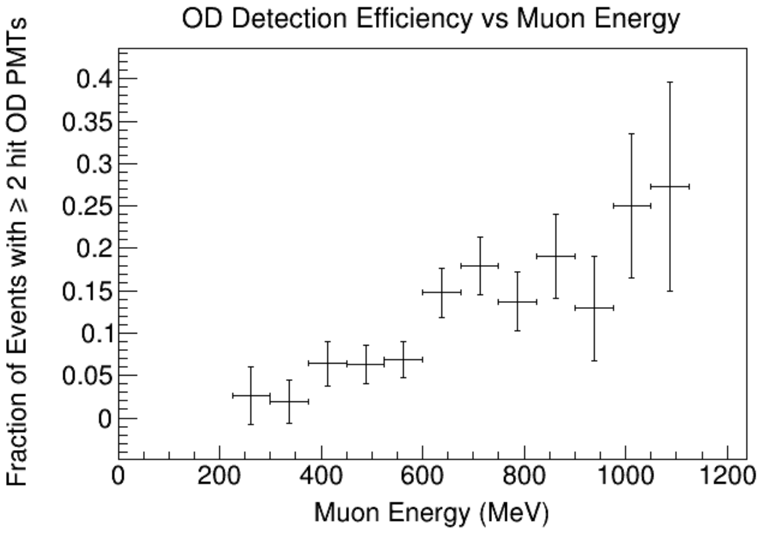

Figure 2.15: The results of the efficiency study for a beam-representative distribution of dirt muons. Efficiency of detection rises with muon energy, but requiring a coincidence of at least 2 hit PMTs resulted in just 30% efficiency at peak energies, falling to around 5% at the peak of the incident muon energy distribution at ∼400 MeV.

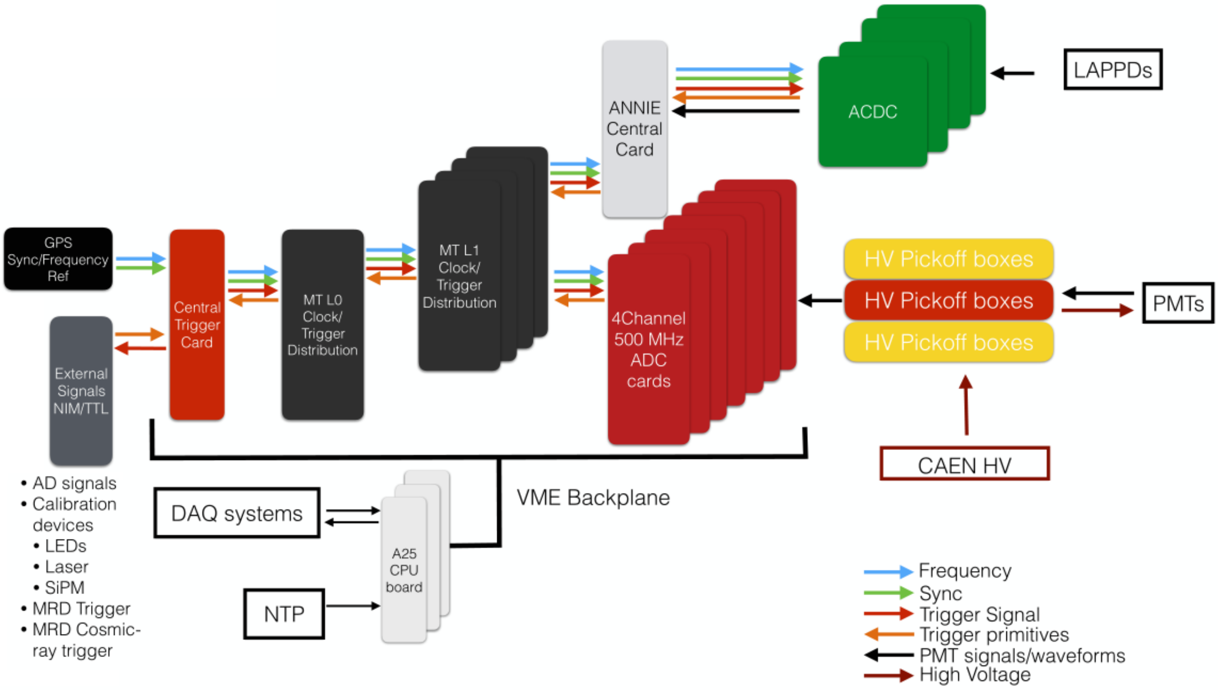

Figure 2.16: A schematic overview of the elements of the DAQ system. Image credit J. Eisch.

2.5 Data Acquisition System (DAQ)

Waveforms from the tank PMTs are recorded by a set of VME digitizer cards originally developed for the K0TO experiment[138]. Waveforms are sampled at 500 MS/s with buffer storage for up to 80µs worth of data on each 4-channel ADC card. Up to eight ADC cards may connect to one Master Trigger (MT) card, which provides clock synchronization and 2-way communication for triggering. Multiple MT cards may then be combined into a higher tier MT card to create a hierarchy spanning the entire PMT array. The top tier cards connect to a Central Trigger Card (CTC) that manages global trigger decision and issuing. This setup is shown in Figure 2.16, along with additional pathways for GPS clock disciplining, external trigger input, CPU data readout, high voltage and LAPPD communication.

In normal beam running a 2 µs prompt window is always recorded at the start of each beam spill, triggered by a signal from the accelerator. During this window each ADC card monitors its PMTs for pulses exceeding a configurable threshold. Each card then communicates the number of channels that exceeded this threshold to its MT card, which propagates the information to the CTC. If multiple PMTs saw a pulse within the readout window the event is considered of interest, and a secondary trigger is relayed back down to the ADC cards to continue readout for

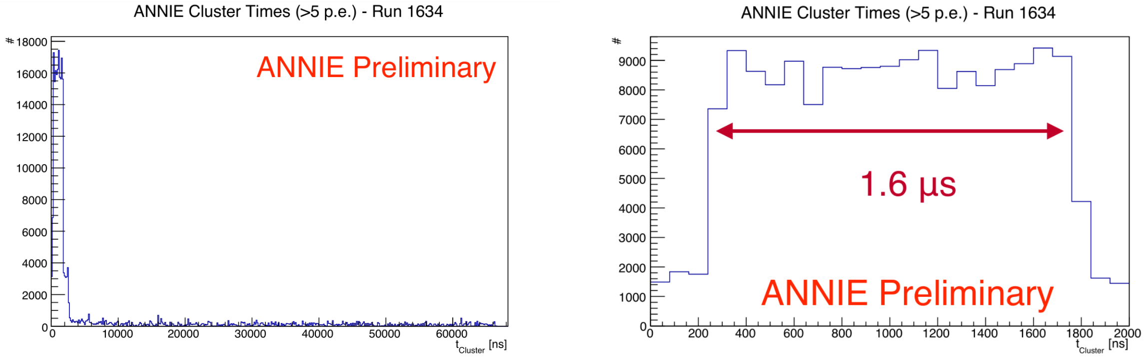

Figure 2.17: Left: A histogram of 'PMT clusters' (>10 hit PMTs within 50 ns) binned in time since the acquisition window start. The acquisition window is triggered by a signal from the accelerator and includes a small pre-beam window. Right: A zoom in on the prompt 2µs region containing the beam spill. The excess demonstrates clear observation of prompt events from the beam.

a further 68 µs. This readout scheme is illustrated in Figure 2.17. The two figures represent the accumulation of all readouts over a particular run. The x-axis shows time since the start of the acquisition window, while the y-axis shows the number of 'hit clusters' (representing at least 10 hit PMTs within a 50 ns window) observed within that time bin. The acquisition window opens just slightly before the neutrinos arrive (triggered by a signal from the accelerator), so the first few bins represent a pre-beam region used for evaluating PMT waveform baselines. Following this is a raised 1.6 µs region demonstrating a clear increase in the number of clusters occurring during the beam spill; a zoom in on this region is shown on the right. Following this is a 68 µs after-beam region in which ANNIE will search for neutron captures. This trigger scheme is flexible and may be changed as data taking progresses.

Along with the nominal beam trigger there are several calibration triggers also in place. The CTC may prime the DAQ to record the next self-trigger event under a number of different configurations, for example to record cosmic ray events or neutron calibration events. It also manages the issuing of triggers to external hardware when relevant, for example to orchestrate the firing of calibration LEDs in time with a 2 µs readout.

Signals from the MRD and forward veto are discriminated by 16-channel LeCroy 4413 CAMAC discriminators with thresholds configurable on a per-card basis. The thresholds have been tuned based on rate measurements to be just above the noise floor in order to maximize sensitivity. The discriminated signals are read out by LeCroy 3377 TDCs configured to record up to 15 hits on each channel over a 4µs window, with a 4 ns resolution. Readout is triggered by the CTC in synchrony with the rest of the DAQ.

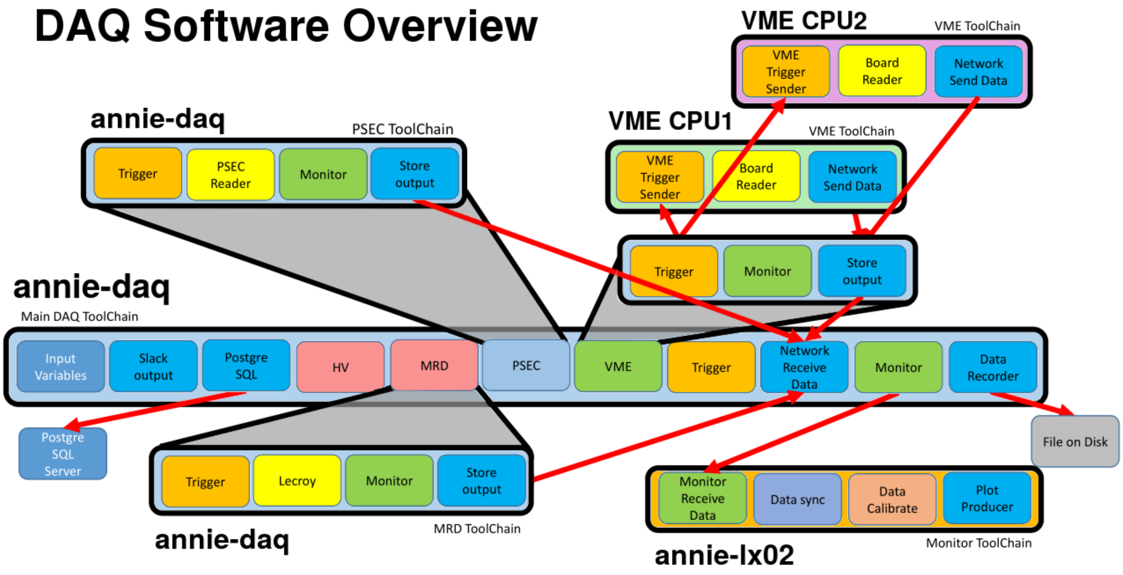

Figure 2.18: An overview of the toolchains controlling the ANNIE DAQ system. Image adapted from one by B. Richards.

2.6 Control and Readout

Central control and readout of all aspects of the DAQ are performed by the ToolDAQ framework. ToolDAQ[139] is a modular, distributable and scalable framework primarily developed for the Hyper-Kamiokande and ANNIE detectors[140]. The system is based around toolchains, comprising a series of tools each performing a specific task. A universal data store provides the means of retrieving data from upstream tools and passing it to downstream tools. Each tool is based on an abstract base class and must provide at minimum `Initialise`, `Execute` and `Finalise` functions. The toolchain initialises by calling each tool's `Initialise` function, then repeatedly loops over each tool's `Execute` function until the toolchain completes, at which point the `Finalise` methods perform clean-up. Toolchains may be nested or distributed across multiple nodes on a network. Separate toolchains provide readout and control of the TDC, ADC and PSEC systems, as well as analysing a subset of data for online monitoring. An overview of the control scheme is shown in Figure 2.18. Various elements run on the main `annie-daq` server, the VME CPUs, and a secondary monitoring server `annie-lx02`.

2.7 High Voltage Supplies

As with many elements of the detector the high voltage supply system is built from components obtained from Fermilab's Physics Resource Equipment Pool (PREP), a stock of equipment donated by past experiments available to users of the facility. High voltage for the tank PMTs is provided by a CAEN SY527 crate equipped with ten sixteen-channel A734P cards. High voltages for the forward veto and MRD are provided by two LeCroy 1440 crates stocked with sixteen LRS 1443 sixteen-channel cards each. High voltage for the LAPPDs is generated by dedicated step-up transformer cards integrated into the electronics in the LAPPD housing, so are treated separately.

Both CAEN and LeCroy crates are continuously monitored over RS232 links by a LabVIEW based slow control system. Each channel is checked roughly once per minute, comparing reported voltage demands against an on-file reference and checking monitored voltage and current draw for compliance with specified tolerances.[1] The most recent readings, graphed histories and logs of past events are displayed on a GUI accessible locally or remotely via VNC (Figure 2.19), and are also published to a simple html website for easy access. Any indication of an issue with the high voltage system triggers an alert on the GUI, the website, and the collaboration Slack[141] workspace. All readings are also logged to file, as are all system changes and alarms. The software provides control of all available parameters, as well as extended functionality such as event scheduling for scenarios such as high voltage scans or scheduled power-downs. The system also communicates with the DAQ by network messaging, allowing high voltage information to be logged as part of the run information included in ANNIE data files.

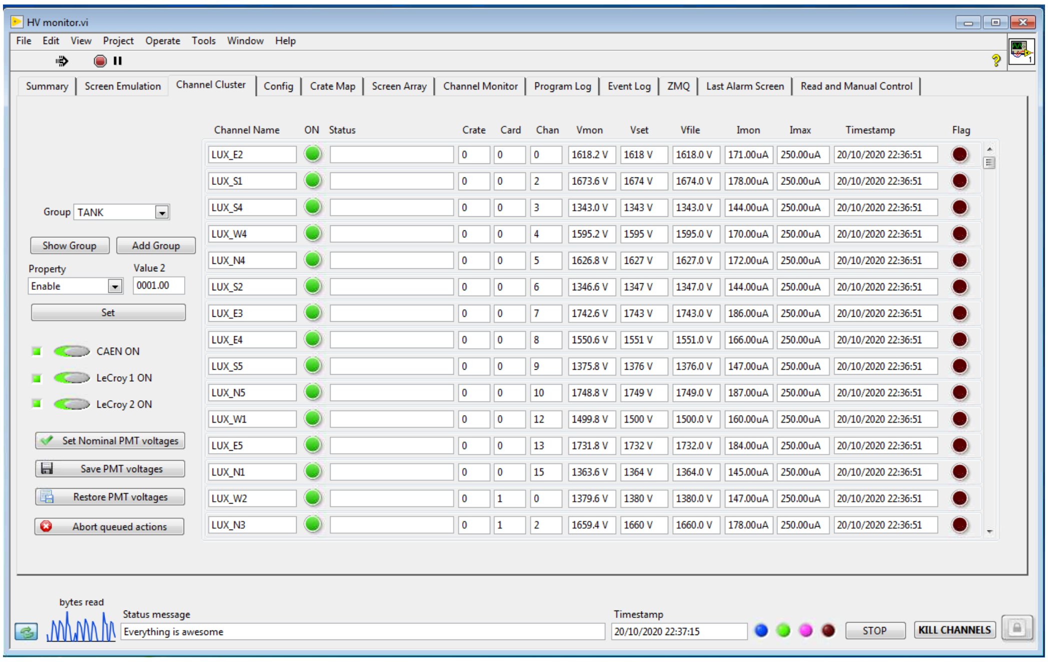

Group TANK

Property Enable

Value 2 0001.00

Channel Name	ON	Status	Crate	Card	Chan	Vmon	Vset	Vfile	Imon	Imax	Timestamp	Flag
LUX_E2	●		0	0	0	1618.2 V	1618 V	1618.0 V	171.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_S1	●		0	0	2	1673.6 V	1674 V	1674.0 V	178.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_S4	●		0	0	3	1343.0 V	1343 V	1343.0 V	144.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_W4	●		0	0	4	1595.2 V	1595 V	1595.0 V	170.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_N4	●		0	0	5	1626.8 V	1627 V	1627.0 V	172.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_S2	●		0	0	6	1346.6 V	1347 V	1347.0 V	144.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_E3	●		0	0	7	1742.6 V	1743 V	1743.0 V	186.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_E4	●		0	0	8	1550.6 V	1551 V	1551.0 V	166.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_S5	●		0	0	9	1375.8 V	1376 V	1376.0 V	147.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_N5	●		0	0	10	1748.8 V	1749 V	1749.0 V	187.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_W1	●		0	0	12	1499.8 V	1500 V	1500.0 V	160.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_E5	●		0	0	13	1731.8 V	1732 V	1732.0 V	184.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_N1	●		0	0	15	1363.6 V	1364 V	1364.0 V	145.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_W2	●		0	1	0	1379.6 V	1380 V	1380.0 V	147.00uA	250.00uA	20/10/2020 22:36:51	●
LUX_N3	●		0	1	2	1659.4 V	1660 V	1660.0 V	178.00uA	250.00uA	20/10/2020 22:36:51	●

Status message Everything is awesome

Timestamp 20/10/2020 22:37:15

Figure 2.19: Screenshot of the channel cluster tab from the LabVIEW GUI for high voltage control.

---

[1] Current draw is monitored for the tank PMTs only as the LeCroys do not provide this functionality.

2.8 Gadolinium Doping

In pure water neutrons lose energy by scattering elastically off nuclei before eventually capturing on a nucleus. As the capture cross-section only grows large at thermal energies most neutrons will thermalise and undergo significant diffusion before being captured. Scattering off hydrogen results in a large fractional energy loss owing to the similar masses involved, so neutrons with an initial energy of $\sim$10–40 MeV, typical of inverse beta decay, can be brought down to thermal energies of 0.025 eV in as few as 20 scatterings. Higher energy neutrons take slightly longer to thermalise, but the correlation is weak as the initial rate of energy loss is higher. The thermalisation time is therefore small in comparison to the diffusion time, with thermalisation taking around $\sim$10 µs compared to $\sim$200 µs of diffusion before the neutron is finally captured[142]. The capture cross-section for thermal neutrons on hydrogen is 0.33 barns, while that of oxygen is just 0.19 millibarns, so essentially all neutrons capture via the reaction $n + p \rightarrow d + \gamma$. The gamma produced is 2.2MeV in energy, which is only just sufficient to Compton scatter electrons above the Cherenkov threshold, resulting in a small amount of detectable light. A study by Super-Kamiokande using a specialised offline search for low-energy activity (neutron capture) correlated with a high luminosity event (positron emission) achieved a 20% efficiency for the detection of neutron captures from IBD[142].

In order to improve the detection efficiency the addition of gadolinium was proposed by J. Beacom and M. Vagins[143]. Gadolinium is a rare-earth lanthanide, but one of the more abundant rare-earth elements. Natural gadolinium has a neutron capture cross-section of 49,000 barns, dominated by the presence of isotopes ^{157}Gd and ^{155}Gd. These have neutron capture cross-sections of 254,000 barns and 61,000 barns, and natural abundances of 14.5% and 15.7% respectively. The capture cross-sections for these isotopes, along with hydrogen, are given as a function of energy in Figure 2.20. The gadolinium cross-section exhibits a forest of resonances in the thermal and epithermal regions where there is a high density of nuclear states that will readily accept neutrons of these energies. More information on the mechanism and measurement of these resonances, and gadolinium capture cross-sections in general, can be found in references such as [144], [145], [146] and [147]. Owing to its drastically larger cross-section only a small amount of gadolinium is required to capture the majority of neutrons, as illustrated in Figure 2.21. A concentration of just 0.1% gadolinium by weight gives 90% efficiency for capturing thermal neutrons and will be used by the ANNIE experiment (the same concentration is planned by the Super-Kamiokande Gadolinium project[88]). Given the 26 tonnes of water in the ANNIE water tank this amounts to the addition of 52 kg of the compound gadolinium sulphate. While

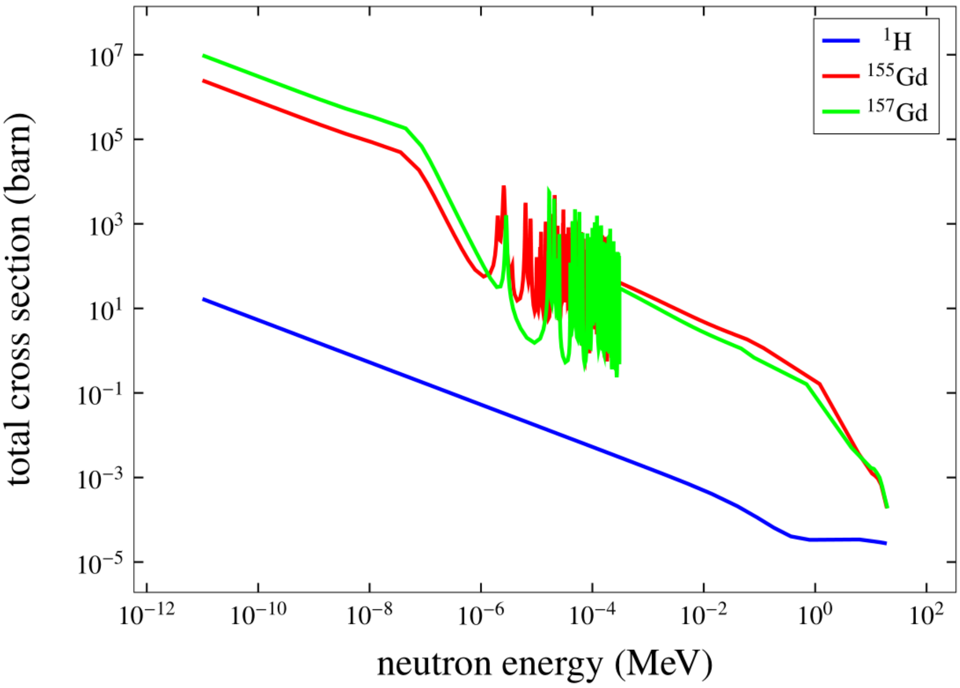

Figure 2.20: Neutron capture cross-sections for the dominant gadolinium isotopes, and hydrogen. Data from the ENDF/B-VIII.0 database[148]

the going price of rare earths tends to be quite variable, this quantity of gadolinium sulphate would be expected to cost on the order of a few thousand pounds sterling; ANNIE was able to acquire the required amount from collaborators also on the Super-Kamiokande experiment.

After excitation by a neutron capture the 157 and 155 isotopes decay via a complex cascade of energy levels[145], releasing a total energy of 7.9 MeV and 8.5 MeV respectively. While this energy is divided between multiple gamma rays, not all of which are energetic enough to result in Cherenkov light, around 92% of gammas are expected to have more than 3 MeV of energy, resulting in a substantially improved detection efficiency. The spectrum of gamma ray energies is shown in Figure 2.22.

As well as increasing the visible energy the addition of gadolinium also significantly reduces the time and distance of diffusion, from ∼200µs and ∼200 cm in pure water to ∼20µs and ∼50 cm with 0.1% gadolinium doping[150]. This permits the use of tighter event selection cuts and a reduction in coincidental backgrounds, and is particularly important in ANNIE as it substantially increases the fraction of neutrons that capture within the fiducial volume.

The ability to tag neutrons in water Cherenkov detectors has numerous physics applications,

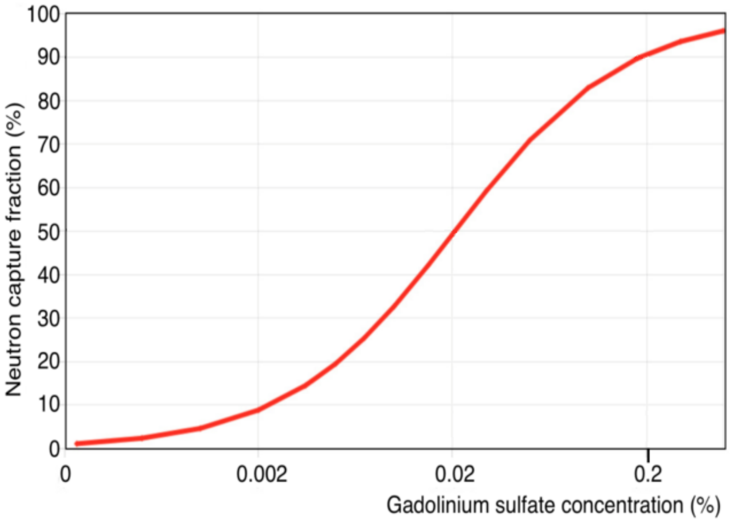

Figure 2.21: Fraction of neutrons that capture on gadolinium as a function of concentration. Figure from [149]

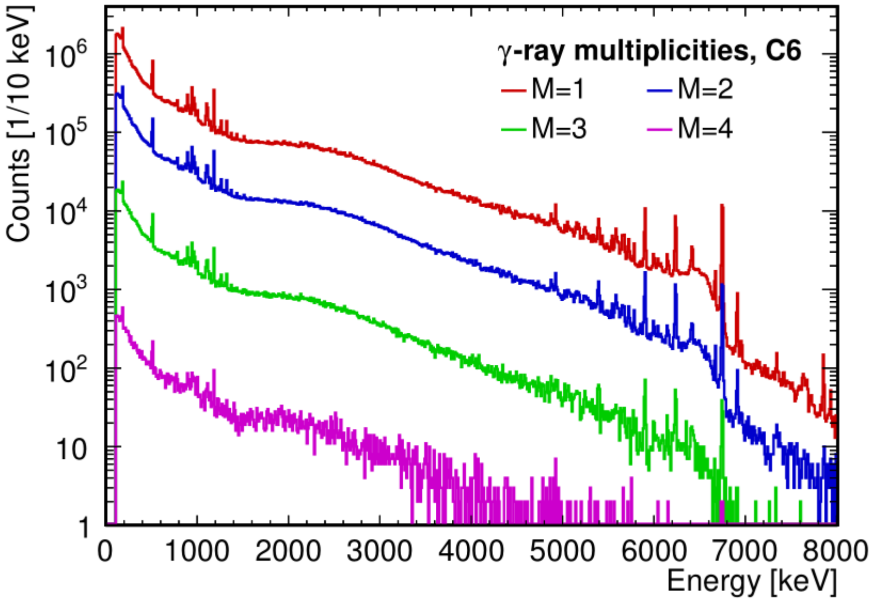

Figure 2.22: Energy spectrum of gamma rays from neutron captures on gadolinium 157. Data from the ANNRI experiment[145]. The curves show events with different multiplicities, based on the number of isolated energy depositions observed in a segmented germanium crystal.

including rejecting backgrounds in proton decay[87] and supernova relic neutrino searches[151], discrimination between neutral current backgrounds and oscillated ν_e interactions[152], and the selection of a purer sample of quasi-elastic interactions for improved energy reconstruction[153], among many others[154]. While the addition of gadolinium can significantly improve the ability to detect neutrons, in order to make quantitative statements about the purity of the resulting samples it is essential to know not just the efficiency of neutron detection, but also the number of neutrons expected in a given class of interaction. This is the information that ANNIE hopes to provide.

2.9 Chemical Compatibility

Free gadolinium ions (Gd^3+) are insoluble in water at normal pH, but gadolinium salts such as gadolinium chloride (Gd_2Cl_3), gadolinium nitrate ($Gd(NO_3)_3$) and gadolinium sulphate ($Gd_2(SO_4)_3$) will readily dissolve in water. Gadolinium nitrate is strongly absorbing within the key region of overlap between Cherenkov emission and PMT sensitivity (recall figure Figure 2.7), ruling it out as a candidate. Gadolinium chloride has a tendency to lift rust, accelerating the rate of any corrosion and potentially harming transparency. This leaves gadolinium sulphate as the preferred compound. ANNIE is using 52 kg of high purity gadolinium sulphate to achieve a 0.2% $Gd_2(SO_4)_3$ concentration, equivalent to 0.1% gadolinium by mass. To verify its purity the absorption spectrum of a 0.2% gadolinium solution was measured between 210 nm and 400 nm. This is shown in Figure 2.23 as the '100% nom. Gd water' curve, and demonstrates that transparency above 300 nm is maintained with only a narrow absorption peak at 270 nm.

While gadolinium sulphate is less reactive than gadolinium chloride it is still slightly acidic, and given the high degree of water purity required it is crucial to ensure this does not cause plastics to leach chemicals or metal components to rust. A great deal of care was therefore taken to ensure maximal rust proofing of any metals placed in the tank, together with the compatibility of all plastics used. The main structure is built from grade 304 stainless steel, while all smaller metal components (nuts, bolts, springs used in PMT mounting structures, stainless steel ties used for cable routing, etc) are either 304 or 302 stainless steel. The 304/302 designation refers to the composition of the steel formulation, with grade 302 containing 17-19% chromium, 8-10% nickel and up to 0.15% carbon, while grade 304 contains 18-20% chromium, 8-10.5% nickel and up to 0.08% carbon. The chromium and nickel are responsible for forming the passive oxide layers that protect against rust, while the additional carbon content of grade 302 helps to increase strength. This additional carbon also increases the risk of chromium carbide precipitation during welding; a process in which carbon and chromium in the steel bond, locally depleting the steel of chromium and compromising the protective oxide layers. Grade 304 steel was therefore selected for the main body of the PMT support structure.

Further protection can be gained by passivation. Passivation is a chemical process that uses an oxidizing acid to remove free iron from the surface of the steel (which can be imparted during processing), stimulating the formation of a complete, uniform chromium/nickel oxide layer. All steel elements used in the PMT support system were passivated by EMI Electro-Max Inc[155] before welding and construction.

Even with prior passivation and best-practice techniques weld joints are still susceptible to corrosion[156]. In order to restore the protective passive layer the weld may be pickled or

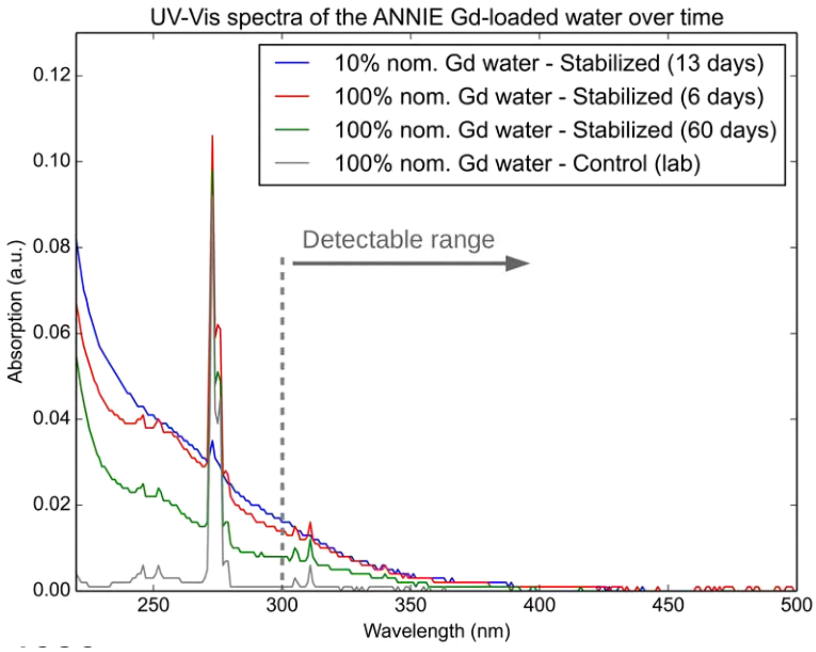

Figure 2.23: Absorption spectra of a control sample of gadolinium loaded deionised lab water and several samples taken from the ANNIE tank. Percentages are concentrations of $Gd_2(SO_4)_3$ relative to the nominal ANNIE loading of 0.2% by mass. The order of samples are 10% nominal loading, then 100% nominal loading after 6 and 60 days respectively. The absorbances fall over time as impurities introduced during loading are removed by the water purification system. Figure by Vincent Fischer.

treated electrochemically. Pickling uses a highly concentrated nitric or hydrofluoric acid to remove the top layer of the steel, allowing a fresh passive layer to reform. The use of hazardous chemicals, however, makes this process potentially dangerous and bad for the environment. Electrochemical treatment (also known as electrolytic cleaning) uses a much milder phosphoric acid in combination with an electrical current to accelerate an electrolysis reaction that detaches contaminants from the steel surface and deposits them on special carbon brush fibres; this restores the surface finish and passivation while being safer and more environmentally sound. ANNIE used a Cougartron CGT-550 electrochemical cleaner to treat all welds on the PMT mounting structure.

As a final precaution the entire inner structure was wrapped in 2-inch wide Teflon tape to slow the leaching of any potential rust into the water. Simple tests of regular (non-stainless) steel left to soak in water showed a stark difference in water quality when wrapped in Teflon tape, as shown in Figure 2.24.

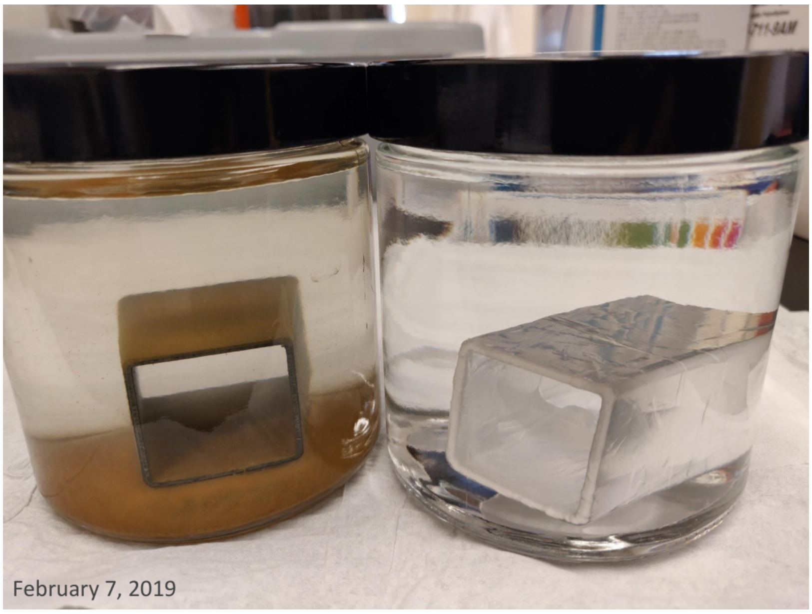

Figure 2.24: Regular steel left in de-ionised water demonstrated that wrapping with Teflon tape significantly reduced the rate of rust leaching into surrounding water. Figure from Julie He.

In order to evaluate the efficacy of each of these treatments, as well as the overall compatibility of all materials that went into the water tank, samples of all bulk materials, treated and untreated, as well as all small parts (nuts, bolts etc.) were soak tested before being released for installation. Soak tests used a 1% concentration of $Gd_2(SO_4)_3$ (five times that used in the detector) in order to accelerate the appearance of any degradation. For samples of the main structure the solution was also heated to 40°C, creating conditions expected to accelerate the rate of any corrosion by 20 times. Throughout the tests water samples were routinely tested for transparency between 200 and 400 nm.
These tests were also carried out on all plastic components, helping to select or rule out various candidate materials. This included elements such as the liner and optical isolation sheet, PMT and calibration LED cable jackets, heat-shrink tubing used in cable extensions and polypropylene piping used for providing a continuous flow of nitrogen through the water. A summary of the results from this testing can be found in [157].
The PMTs and their mounting assemblies were also tested for chemical compatibility. There are three different mounting types, shown in Figure 2.25. The ETEL and LUX tubes, shown in Figure 2.25a, are both waterproofed by enclosing the back end of the PMT within a cylindrical PVC section, capped with an acrylic window and bonded to the PMT bulb with epoxy. The PMTs are fixed to the inner structure by a set of wings glued to the cylindrical body. The WATCHBOY, WATCHMAN and new ANNIE PMTs are waterproofed by Hamamatsu via en-

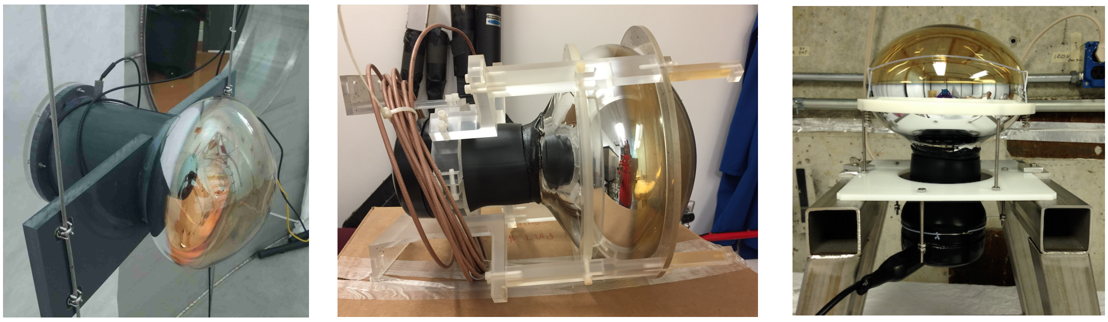

Figure 2.25: Left: A LUX PMT in its original experiment; the holders are unchanged for use in ANNIE. The same style of holder is used for the ETEL tubes. Middle: A PMT from the WATCHBOY experiment in its holder. The same style of holder is used for the WATCHMAN tubes. Right: A new PMT purchased by ANNIE, in a holder of the ANNIE design. The mounting of each of the styles on the support structure can be seen in Figure 2.8.

capsulation in potting compound. The WATCHBOY and WATCHMAN PMTs mount onto the inner structure with an acrylic holder (Figure 2.25b), while the new ANNIE PMTs are held by a PVC plate and ring, together with a stainless steel wire hoop (Figure 2.25c). Sample LUX and WATCHMAN PMTs, including holders, were soak tested for over a year in 1% gadolinium sulphate solution, and neither showed any signs of degradation.

Figure 2.23 shows the absorption spectra of water samples taken from the tank during an early 10%-of-nominal gadolinium loading phase, and then at various times following the full loading. The curves show no sign of corrosion or leaching of harmful contaminants into the ANNIE water. The absorption seen below 300 nm is not a problem, since it lies outside the range of sensitivity of the PMTs, but is monitored closely as an early indication of potential problems. At the time of writing, 11 months after the initial loading, there are no signs of any deterioration.

2.10 Water Quality

Despite the increased energy released by captures on gadolinium only $\mathcal{O}(1000)$ photons are generated by a neutron capture. Accounting for 8% photocoverage and 30% quantum efficiency this amounts to just ∼20–30 photo-electrons. Any impurities in the water will absorb and scatter light, reducing this number even further and compromising the ability to reconstruct the event, so the water must be continuously cleaned by in-situ filtering. In pure water this is not a significant challenge; a long history of experiments have achieved close to perfect transparency by a series of filters, deionising resins and reverse osmosis membranes. Although many of these components are still suitable for use with gadolinium loaded water, typical deionisers and reverse osmosis filters would remove the Gd^{3+} and SO_4^{2-} ions that are now desired. This problem was first addressed by the EGADS project ('Evaluating Gadolinium's Action on Detector Systems')[158], which demonstrated a molecular bandpass filter that selectively extracts the desired ions from the water, allowing all remaining impurities to be filtered out before recombination[151]. While this technique works very well it is expensive and consumes considerable space and power.
To address this issue ANNIE collaborators at the University of California in Davis developed an ion exchange resin capable of filtering out all ions *except* Gd^{3+} and SO_4^{2-}[159]. The relaxed water transparency required by the smaller ANNIE tank also permits bypassing the reverse osmosis stage, resulting in a far smaller and cheaper system that maintains more than adequate water transparency with very little loss of gadolinium. Conservative estimates based on small scale tests indicated that 2 years of continuous filtration would result in a ∼10% (∼5 kg) loss. The gadolinium levels of ANNIE will be routinely monitored through a combination of conductivity measurements, transparency measurements, neutron source calibration runs and spectroscopy; maintaining a consistent concentration is expected to be readily manageable.

A schematic of the water filtration system is shown in Figure 2.26. The various stages include:

1. An initial 5 μm filter to remove dust and large particulates;
2. A 254 nm UV lamp to kill bacteria;
3. A Total Organic Carbon (TOC) lamp, which emits 185 nm UV light tuned to break the carbon bonds in organic molecules;
4. A specially formulated deionizing resin, which captures the radicals produced by the TOC lamp along with any other ionic impurities, while allowing Gd^{3+} and SO_4^{2-} through;

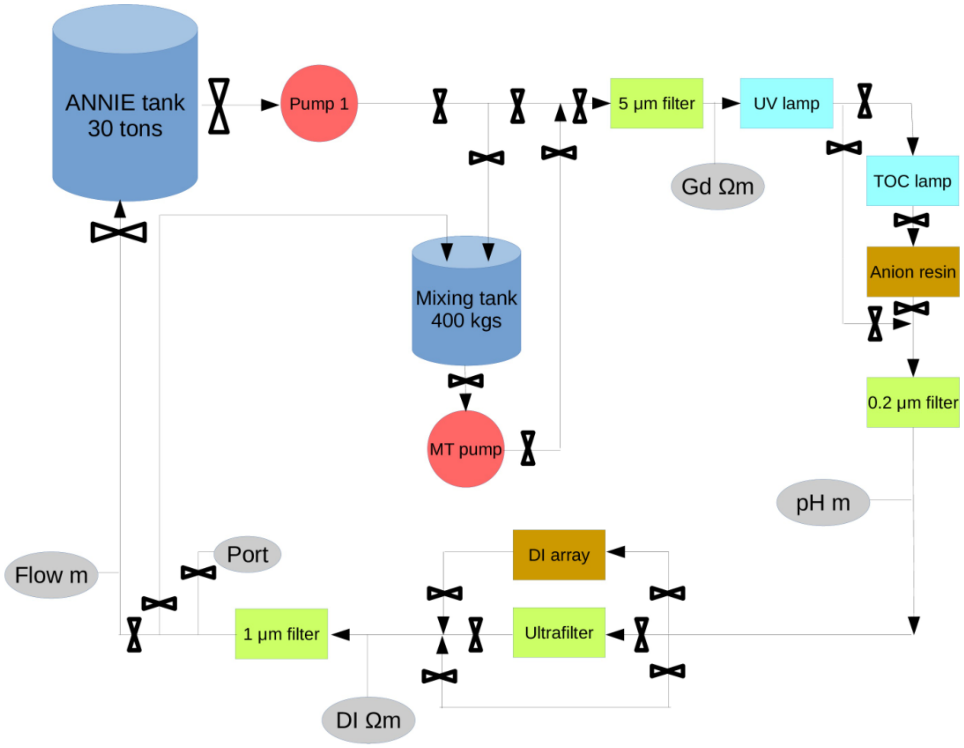

Figure 2.26: Schematic of the ANNIE water filtration system. Figure credit V. Fischer

5. A 0.2 µm filter to remove finer particulates;
6. A normal deioniser resin, used during initial pure water fill but bypassed for gadolinium-loaded operation;
7. An ultrafilter featuring a PVDF membrane with 30 nm pores. At the expected pH of 5-6 iron forms large molecular colloids that will be filtered out, while gadolinium will be able to pass through;
8. A final 1 µm filter, used to catch any possible resin beads when bypassing earlier parts of the system during the initial loading or routine maintenance;

Along with the water filtration a continuous supply of nitrogen is bubbled in through the bottom of the tank and a blanket of nitrogen is maintained over the top of the tank, which helps to lower the oxygenation level of the water and suppress the growth of biologics. The experimental hall is also kept air conditioned to a maximum of 20°C, which further helps to suppress bacterial growth, and also aids with the dissolution of gadolinium.

During each stage of loading a 1% gadolinium solution (half the solubility limit) was prepared in a 110 gallon mixing tank and run through the filtration system to remove any potential impurities before being injected into the main detector tank. Detector loading was performed in three stages. An initial stage with 1% of nominal concentration (0.002% $Gd_2(SO_4)_3$) was carefully monitored for 2 months for any signs of degradation. With no detrimental effects seen a second stage of 10% of nominal concentration was then monitored for a further 2 weeks, before finally ramping up to 100% of nominal concentration over the course of a final week. Throughout the entire loading procedure water samples were taken twice a day and tested for optical transparency, along with continuous monitoring of water conductivity, pH, flow rate and water temperature. Consistently high levels of optical transparency were seen throughout, as shown in Figure 3.8. Under normal operating conditions the water in the tank is expected to circulate through the filtration system once every 2.5 days.

ANNIE will use the Mini Irvine Device for Evaluating Attenuation Length (MIDEAL) for routine monitoring of optical transparency. A schematic representation of the MIDEAL setup is shown in Figure 2.27. The system works as follows: a laser light is shone through a beam-splitter, with one path directly illuminating a photosensor and the other passing through a column of water onto a second photosensor. The amount of water in the column is then automatically varied, via pneumatic valves for fill and drain, between 0 and 3 meters. The relative amount of light on each sensor as a function of column length allows a highly accurate determination of transparency that is insensitive to the absolute photosensor responses. A series of 7 different lasers automatically measures the attenuation at several points across the Cherenkov spectrum; 337nm, 375nm, 405nm, 445nm, 473nm, 532nm and 595nm.

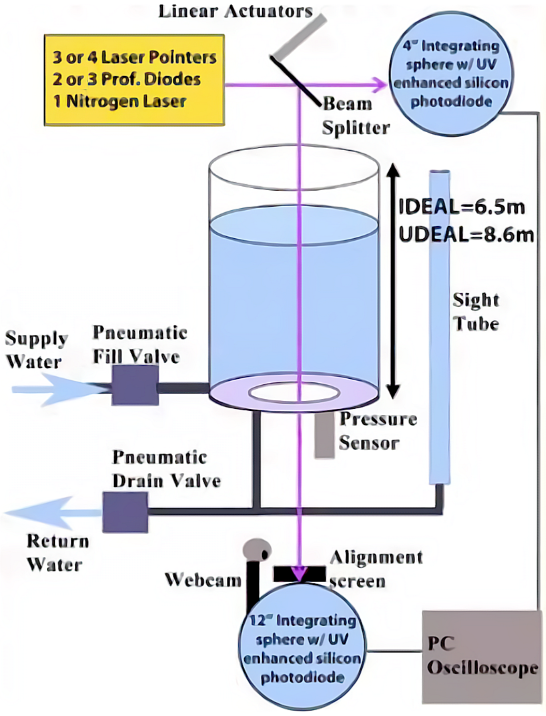

Figure 2.27: Schematic of the MIDEAL automated transparency monitoring system. Figure credit M. Smy

2.11 Large Area Picosecond Photodetectors (LAPPDs)

2.11.1 Motivation

For cross-section measurements the definition of the target volume is a critical parameter that feeds directly into the expected number of events. Minimizing the uncertainty on the target volume means achieving the best resolution on the interaction vertex. As will be seen in section 6.3 the reconstructed vertex is obtained from a likelihood fit to the arrival times of photons at each photosensor, so any uncertainty on the arrival times translates directly into uncertainty in the interaction position. An uncertainty of 3 ns on the arrival time, typical of the transit time spread of a 10-inch PMT, corresponds to an uncertainty of ∼60 cm on the point of photon emission. The possible vertex resolution is therefore ultimately limited by the timing characteristics of the photosensors. In order to achieve a better vertex resolution ANNIE will be pioneering the use of Large Area Picosecond Photodetectors. As the name suggests these have a much better timing resolution, allowing far better vertex reconstruction and fiducialization.

This is not just important for ANNIE's cross-section measurements. Neutrons can diffuse upwards of 50 cm from their origin before being captured, with a greater diffusion length in the beam direction due to the momentum transferred from the neutrino[160]. The distribution of capture distances predicted by Monte Carlo simulations is shown in Figure 2.28. The attenuation length for gamma particles emitted during Gd decay is of a similar scale. If the neutron escapes, or if the gamma rays leave before depositing sufficient visible energy, the neutron will go undetected. The efficiency of detection is therefore lower for interactions occurring close to the tank walls, and for events further downstream. This variation is shown in Figure 2.29, which shows the fraction of final state neutrons producing at least 5 photoelectrons of detected signal as a function of the neutrino interaction position. Accounting for this requires the definition of a fiducial volume to determine the integrated efficiency. Since the efficiency varies quite rapidly over distances as short as 10 cm a vertex resolution at least as good as this is crucial for providing an accurate efficiency correction, and ultimately an accurate measurement of neutron yield.

2.11.2 Development & Construction

The LAPPD Collaboration[161] was formed in 2009, following community recognition of a need for picosecond capable detectors in extracting maximum information from particle collisions. The collaboration focused on the development of Micro Channel Plates (MCPs) - detectors consisting of a photocathode stacked over thin layers of porous substrate held under high voltage

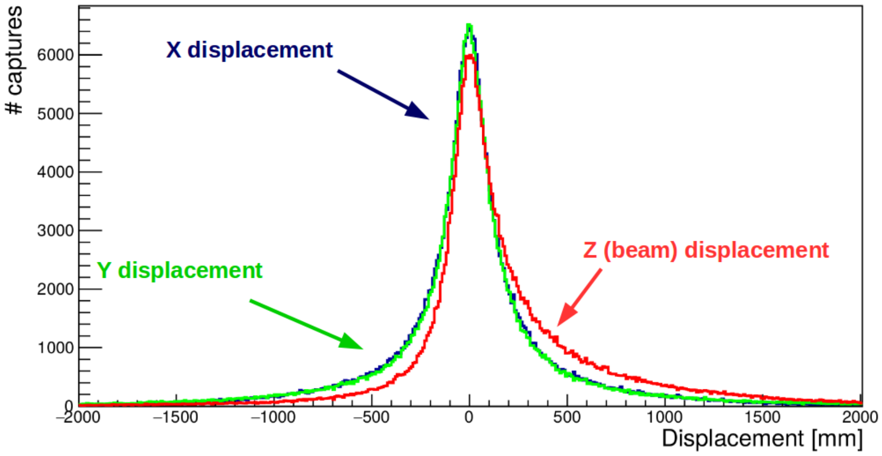

Figure 2.28: Neutrons may diffuse upwards of 50cm from their origin before capturing on a nucleus. Neutrons that escape the tank before during this propagation will go undetected. Plot by Vincent Fischer.

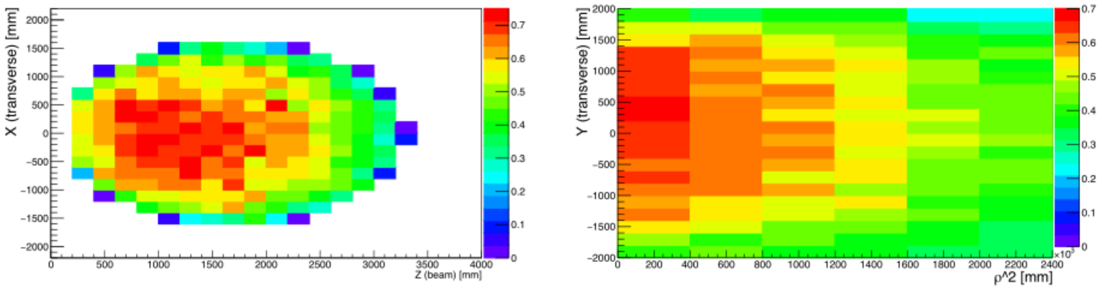

Figure 2.29: A plot of the fraction of final state neutrons whose capture produces at least 5 photoelectrons of detected signal, binned according to the position of the original neutrino interaction (not the neutron capture). The left plot is a top-down view of the tank, integrated over the height within the fiducial volume. The beam enters from the left and exits to the right. The positive skew of neutron travel distances illustrated in Figure 2.28 results in a greater chance of neutrons escaping when the interaction occurs downstream; as a result the region of maximum efficiency is shifted upstream (left in the plot). In the right plot positions are binned in height within the tank Y against radial unit $\rho^2 \equiv x^2 + y^2$. The 2.5 tonne fiducial volume for neutron captures is defined by the high efficiency region (red), but accounting for this efficiency requires a good knowledge of the interaction origin. This makes vertex reconstruction resolution a key parameter for the experiment. Plot by Vincent Fischer.

(Figure 2.30). Electrons ejected from the photocathode cascade through the subtrate pores, scattering off the walls and producing an avalanche. The resulting current pulse is deposited on a series of stripline anodes. Each anode is read out at both ends, with the difference in pulse arrival time giving the hit position along its length. The hit charge will be spread across multiple anodes, so the hit poisition transverse to the striplines can be estimated from the centroid of this spread. The critical fast timing is achieved thanks to their compact design. With only a few centimeters between anode and cathode the electron path lengths are far more constrained, reducing the spread in transit times and producing much narrower pulses than conventional PMTs. This also helps to reduce their susceptibility to magnetic fields[162] which large PMTs tend to suffer from, as it reduces collection efficiency. While they clearly fit many requirements, conventional MCPs are limited by manufacturing processes to a very small size, with the largest available offering an active area of just 5x5cm. They are sufficiently expensive to prohibit their application in the numbers required to achieve high coverage, are slow to manufacture and suffer a relatively short lifespan ($\sim$1C/cm^2) due to a susceptibility to ion feedback. Making MCP detectors meet the needs of next generation experiments meant offering larger areas, better lifetimes, and shorter and cheaper fabrication methods. All of these goals were realised by the LAPPD design, which officially moved from R&D to commercialization in 2012. The interested reader can find a comparison of other fast timing detectors in ref [163].

The construction of an LAPPD is shown in Figure 2.31. The photocathode is thermally evaporated onto a fused silica front window. This is hermetically sealed to a glass body by coating the contact surfaces with NiCr and Cu layers before contact-welding with an indium wire. The body encapsulates a stack of two MCPs, separated by resistive spacers that operate as a high-voltage divider. Individual HV connections are also provided to allow finer control of the potential difference across each stage. Stripline anodes are laid onto a bottom glass grounding plane by silk-screening a silver-loaded paste and firing at high temperature. The anode plane is then frit bonded to the bottom of the glass body[1], with the anodes passing through the frit bond to provide readout access while maintaining the internal vacuum. Finally a metal sheet covering the back side of the glass plate provides the current return path and shielding against EM noise.

Traditional MCPs are formed by inserting low-Z glass rods into lead glass tubes. These are then heated and drawn out to produce thin fibres with a lead glass cladding and a light glass core. Many fibres are bundled together, fused, and drawn out again. The process is repeated

---

[1] Frit bonding involves laying down a fine layer of ground glass (glass frit) between the two layers and firing at moderate temperature until the frit melts, joining the two bodies and providing a strong hermetic seal.

until the core diameters reach the desired size; typically 10-100 μm. Sections of the bundles are then sliced out to form wafers. The core glass is chemically etched away, leaving the hollow pores of the MCP. This is then heated in the presence of hydrogen, resulting in the formation of an alkali surface layer that provides the resistive and emissive properties required for electron amplification.
This process has a number of drawbacks. Since the MCP is only made from one material the resistive and emissive properties must be balanced against the required mechanical properties. The hydrogen firing process makes the MCP brittle, and potassium in the lead glass contributes to high background rates via its radioactive decay. The initial electrical properties of the MCP are not stable, requiring a long ‘burn-in’ time before quasi-stable gain is achieved, although this continues to fall through the lifetime of the device.

To improve on this design LAPPDs use cheaper, stronger, and lower radioactivity borosilicate glass together with Atomic Layer Deposition (ALD) to apply resistive and emissive coatings. This allows independent tuning of the mechanical and electrical properties, via a cheap, scalable, industry-standard batch processing method. The capillary array is built in a similar way to conventional MCPs, but hollow light glass tubes are used rather than filled lead glass ones, so no chemical etching step is required. After drawing down to the desired pore size many capillary arrays are assembled into a large block and fused to produce the required 20 cm-square area. After cutting and polishing the wafers are cleaned in preparation for thin film application. An initial 100 nm resistive layer is applied via ALD to support the constant high voltage across the plate, followed by a 20 nm thick emissive layer that tunes the secondary electron emission properties to provide a high avalanche gain. A final conductive layer, applied by thermal evaporation to the top and bottom of the plate, supports the distribution of high voltage across the active area. The result provides a robust, high gain, low noise MCP with no burn-in required. The tiles have been demonstrated to withstand total extracted charges in excess of $7\,\mathrm{C/cm^2}$ without loss of gain, and a factor of 4 lower dark noise than conventional MCPs. A summary of LAPPD properties is given in Table 2.4.

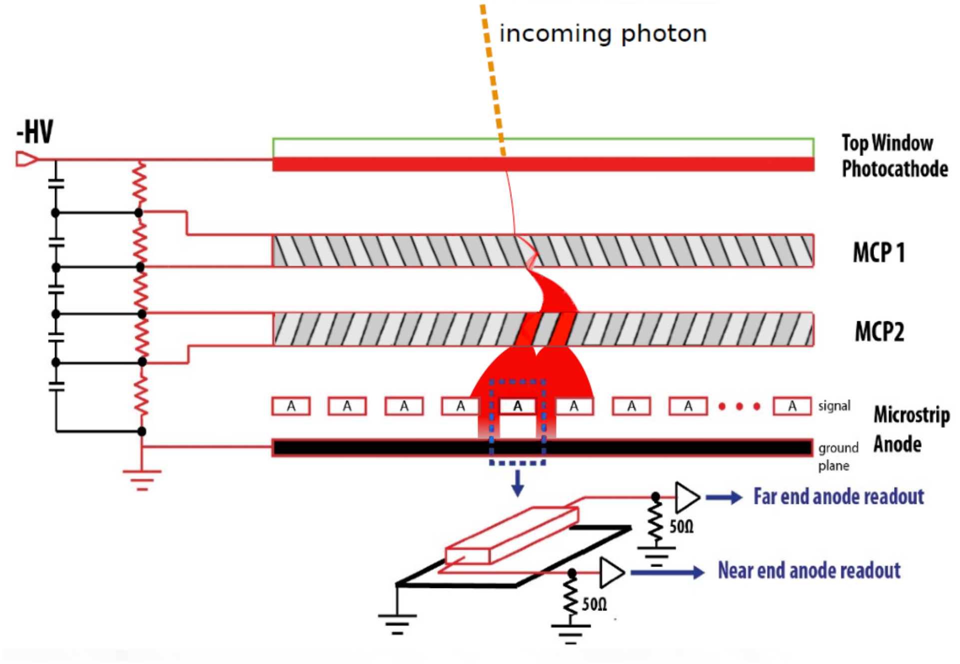

Figure 2.30: Interaction of a photon with the LAPPD, resulting in an avalanche through the microchannels that may spread across multiple anode strips. Figure adapted from [164]

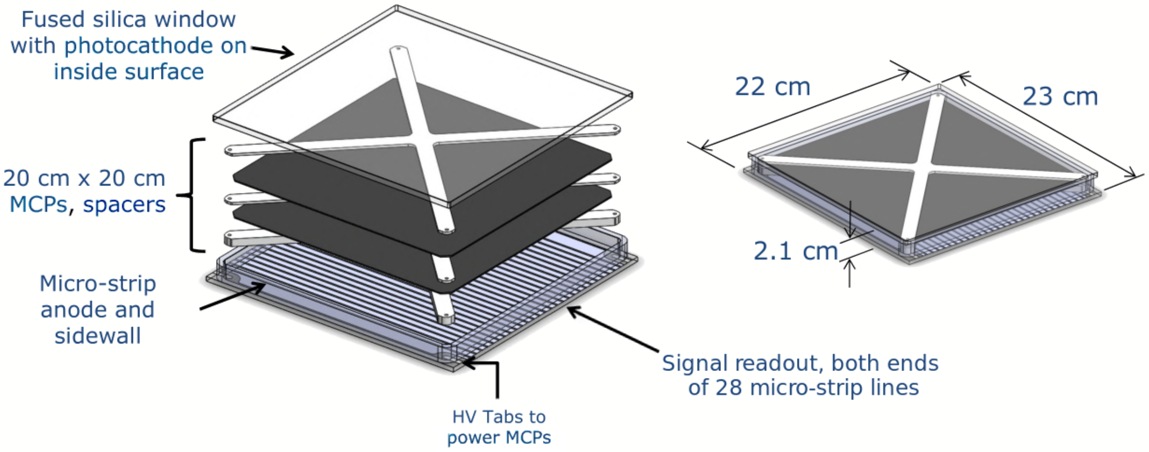

Figure 2.31: Construction of the LAPPD. Figure from [165]

Property	Value	Note
Outer Dimensions	22 x 23 x 2.1 cm	-
Body Material	Borosilicate glass	-
Window Material	Borosilicate glass	-
Photocathode Formula	Multi-alkali K_2NaSb	Low resistance to support high rates
Spectral Response	160-650 nm	-
Quantum Efficiency	20-30% at 365 nm	>35% at 254nm
QE Uniformity	<1.5%	-
MCP Total Area	19.5 x 19.5 cm = 380 cm^2	-
MCP Active Area	350 cm^2 (92%)	After accounting for spacers
MCP Thickness	1.5mm	8° pore angle
Gain	mid-10^6	uniformity <15%
Dark Noise	0.1-1 cnt/s/cm^2	-
Supply Voltage	-2750 V	Photocathode voltage
	450 V	Photocathode to top MCP
	975 V	Across each MCP
	200 V	Between MCPs
	200 V	Bottom MCP to anode
Parallel Resolution	2.4 mm	From time differential at anode ends
Perpendicular Resolution	0.76 mm	From charge centroid on adjacent strips

Table 2.4: Information from [165]

2.11.3 Readout and Electronics

Each LAPPD has 28 anode striplines which may be read out at both ends, or at only one end leaving the other unterminated. When reading out both ends the position along the strip is determined using the difference in arrival time of the pulse at each end. When reading out just one end the time difference between the direct and reflected pulse may be used, halving the required channel count but increasing the complexity of signal deconvolution. ANNIE will be using two-sided readout.

Signals are digitized at 15 GSa/s using a specially developed 6-channel CMOS ASIC[166]. Waveform sampling enables the best time resolution by a fit to the leading edge of the waveform[167], as well as enabling intelligent afterpulse rejection and multi-photon disambiguation[168]. Five digitizing ‘PSEC4’ chips are required to read out each side of the LAPPD, and are provided by one ‘**Ac**quisition and **D**igitization with PSE**C**4’ (ACDC) card. A pickup board mounted to the back of the LAPPD supports two ACDC mezzanine cards and provides the interface to the

LAPPD anodes. Due to issues with the early revision ACDC cards that ANNIE has on-hand, the pickup board also splits the signals out to a separate trigger card that provides threshold detection and an N-channel trigger to the ACDC cards.

Power is provided by a Low Voltage - High Voltage (LV-HV) board, which receives low voltage power and generates the required voltages for all other electronics. This includes five separate high voltage outputs, one for each internal stage the LAPPD, and 3 low voltage outputs for the ACDC and trigger cards. Integrated CANBus communication provides control and monitoring of the tile voltages and currents, as well as monitoring of integrated temperature, humidity and leak sensors. The pickup board, two ACDC cards, trigger card and LV-HV board are all integrated into the ANNIE waterproof housing shown in Figure 2.9.

Slow control is managed by a Raspberry Pi, while waveform data from the ACDC cards is passed into a VME-based ANNIE Central Card (ACC). Each ACC supports up to eight ACDC cards and interfaces to the ANNIE DAQ as shown in Figure 2.16.

Properties of the PSEC4 digitizers are shown in Table 2.5. Each chip provides a buffer depth of just 256 samples per channel, giving a relatively limited readout window of 25.6 ns at 10 GSa/s. Reading out the full buffer on all channels of the chip takes ∼20 µs, so each LAPPD may only be read out once each beam spill. Fortunately the LAPPDs are primarily used in reconstructing the prompt event, so this does not represent a significant drawback. In fact, even during waveform readout it is possible for the ACDC cards to operate in 'TDC mode' whereby hit times are still recorded. This means that even for delayed neutron captures the additional photosensor coverage offered by the LAPPDs will not be lost.

Parameter	**Value**
Channels	6
Sample Rate	4-15 GSa/s
Analog Bandwidth	1.5 GHz
Samples/channel	256
Signal Voltage Range	1 V
DC RMS Dynamic Range	10.5 bits
RMS Voltage Noise	700 µV
ADC Conversion Time	1.4 µs
Readout Time	19.2 µs
Maximum Sustained Trigger Rate	50 kHz

Table 2.5: Properties of the PSEC4 digitizer chips. Properties from [166]

3

Calibration

Cutting edge hardware is of little use if its response is not well understood; each PMT and LAPPD is unique and must be calibrated to take full advantage of its sensitivity. This involves ex situ measurements, to characterise detector components before assembly and determine the optimum layout, and in-situ measurements to assess the response of the detector as a whole and to ensure this remains stable throughout running.

3.1 PMTs

3.1.1 Test Stand

Initial characterisation of the water tank PMTs was carried out with a test stand built from CAMAC and standalone electronics. Custom software automatically carried out tests on up to 4 PMTs at a time. A schematic of the test setup is shown in Figure 3.1.

The first test performed was a cooldown test. Exposure to bright (ambient) light puts the photocathode in an excited state, resulting in an increase in the rate of spurious noise pulses. Once placed in a dark environment the excitation will decay away exponentially, but in some cases this may take many hours. In order to establish a nominal response during characterisation, and to understand the time for the detector to reach steady state after the tank has been opened, each PMT was allowed to reach equilibrium before further tests began. The dark rate was recorded once every minute for two hours, sufficient for all PMTs to reach steady state. A summary of the results is shown in Figure 3.2.

Perhaps the most important characteristic of a PMT is how the gain varies with the applied voltage. Too low a voltage results in low gain and poor single photon detection efficiency; too high results in increased dark noise, higher current draw, increased afterpulsing and even a

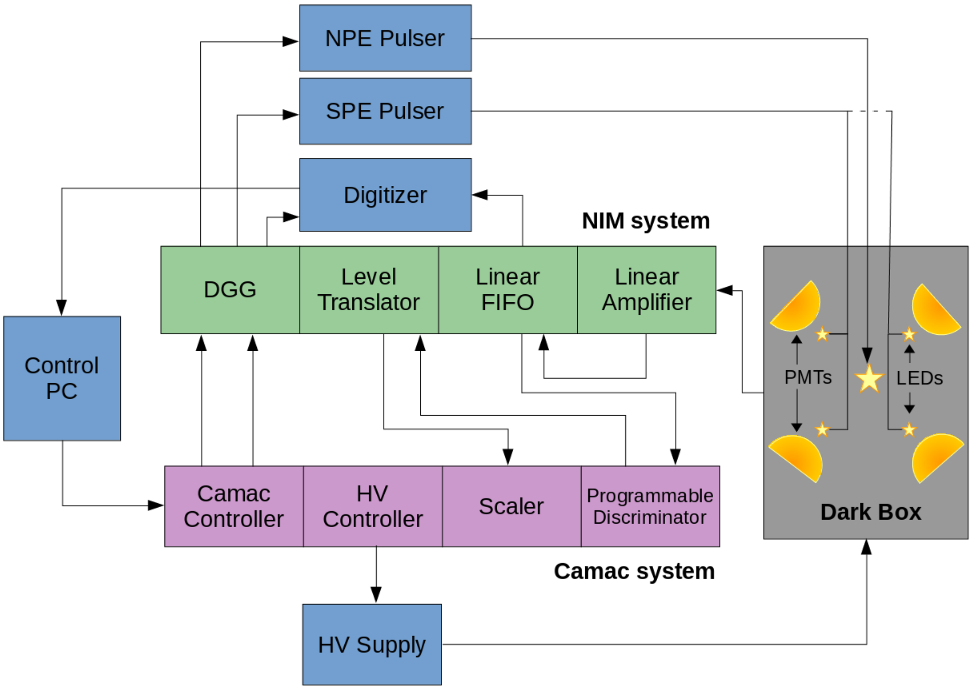

Figure 3.1: Trigger and data acquisition schematic for the PMT characterisation test stand.

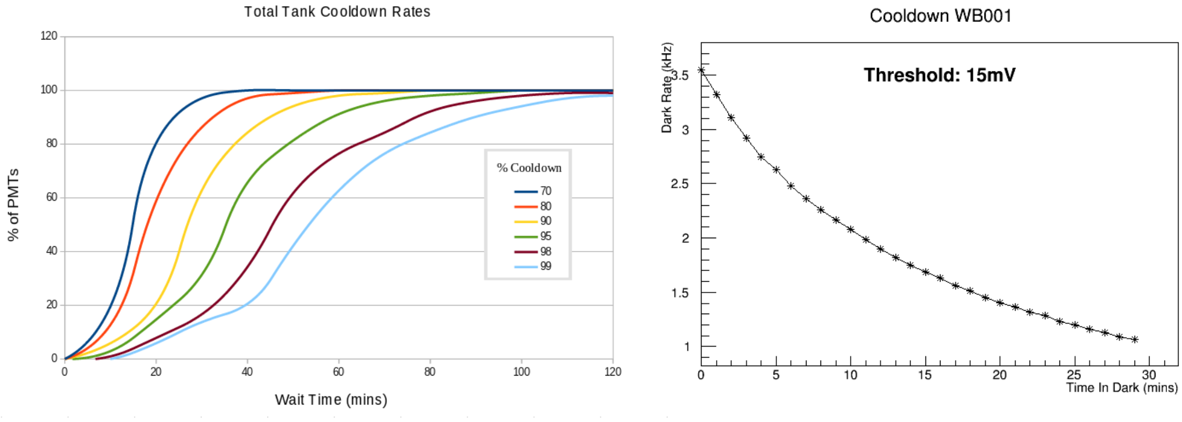

Figure 3.2: Left: Photocathode cooldown times for all tested PMTs. Each colour defines the level of dark rate saturation, while the curves plot the fraction of PMTs that have reach that level of saturation in the given time. Right: An example cooldown curve.

reduced PMT lifetime. On a single photosensor level, the optimum voltage is the minimum that achieves good separation of the single photo-electron peak from thermal and electronics noise. However, often reconstruction algorithms will require uniformity of the photosensor responses, so voltages may also be chosen to normalise the response across a set of PMTs in a detector. It is therefore ideal to have a good understanding of the gain over a wide range of voltages.

In order to evaluate the gain-voltage curve the PMTs were illuminated with very low level light, so that only around 10% of events produced a PMT pulse. At this level the majority of events that do have a pulse will have only seen a single photon, giving a pronounced single photo-electron (SPE) peak with minimal contamination from higher multiplicities. At each voltage 20,000 such events were recorded, with waveforms digitized at 500 MHz and pulse charges calculated using a dynamic integration window. An example of the resulting charge distribution is shown in Figure 3.3. The primary peak at zero charge is the ‘pedestal’ arising from events with no detected photon. The spread of this peak is largely an indication of the system electronics noise. The next dominant peak is the single photo-electron response in which the ejected electron follows the nominal path. The centre of this peak gives the nominal gain of the PMT, and its width includes contributions from system noise and the stochastic nature of the avalanche. Between these two peaks there are a number of other contributions, arising from events in which the primary electron follows an atypical path. This includes events in which the photo-electron misses or initially back-scatters off the first dynode, cases where the photoelectron undergoes sub-optimal acceleration due to electric field inhomogeneities, or even events in which the first electron is ejected directly from the first dynode via the photoelectric effect[169]. If not properly accounted for these contributions can impact the fit to the nominal peak, skewing the measurement of nominal gain. Fits were therefore performed using an algorithm published by the DEAP-3600 collaboration[170], which includes four contributions shown in Figure 3.4; details may be found in the cited reference. An example of the resulting fit to data is shown in Figure 3.3. Fit convergence was found to be particularly sensitive to the initial estimate of the peak position, which was often difficult to extract from the wide variety of distribution shapes. To provide a robust initial estimate the algorithm described in ref [171] was used. The extracted gain-voltage curves are shown in Figure 3.5.

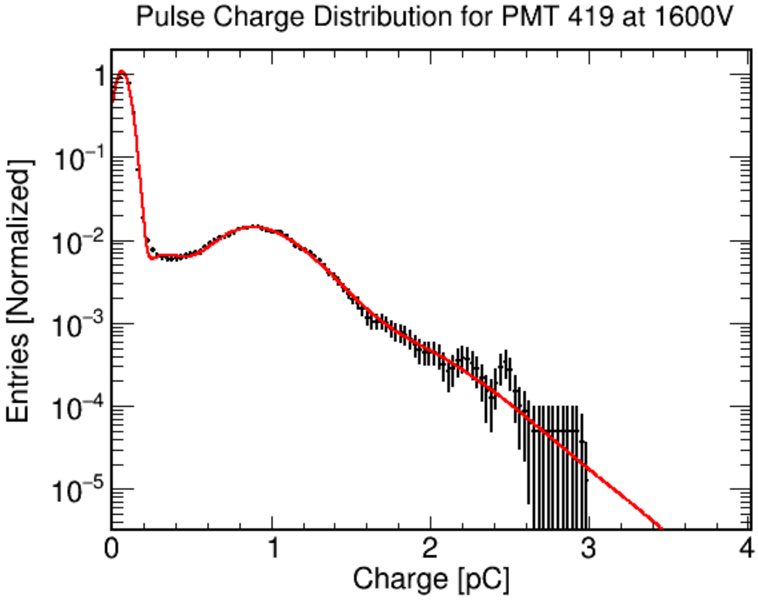

Figure 3.3: An example pulse charge distrbution, together with the fit used to extract the gain

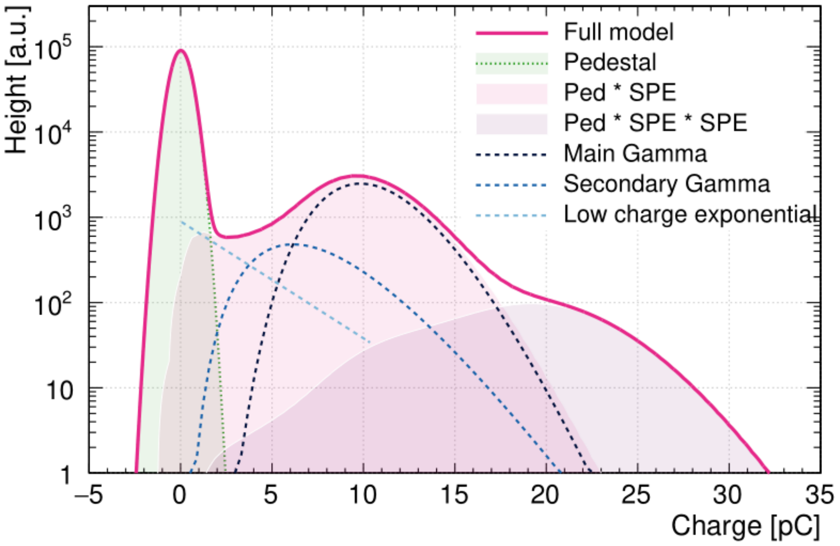

Figure 3.4: Components of the fit algorithm from ref [170]. Figure reproduced from this reference.

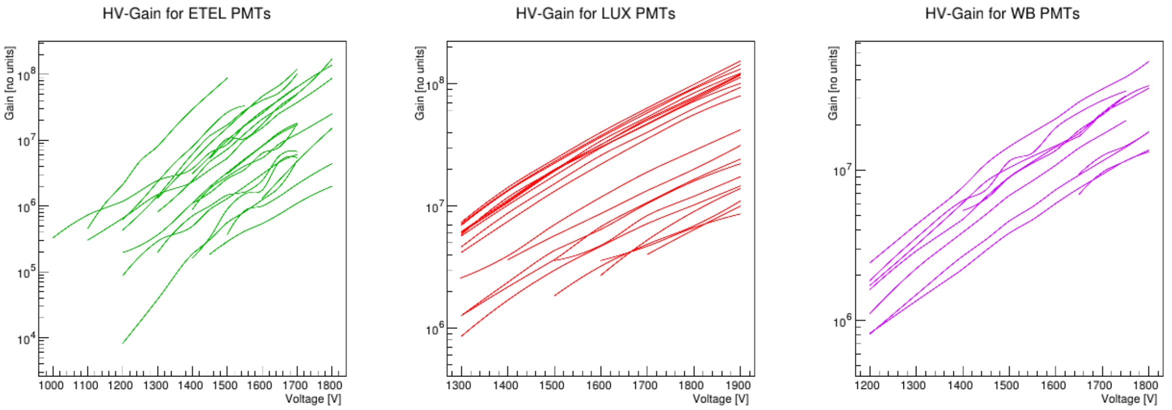

Figure 3.5: Gain variation with voltage, as measured by the PMT test stand, for three sources of ANNIE PMTs. Each line represents a different PMT.

The test stand also performed measurements of afterpulsing; a phenomenon with particular relevance for the neutron yield measurement. Afterpulsing is the occurrence of correlated, delayed pulses following an initial pulse in response to light. It is caused by the ionization of residual gases within the PMT vacuum, generating ions that drift to the photocathode and liberate new electrons on impact. The type of ions and their drift velocities tend to give several characteristic timescales to afterpulsing[172], with delays of the order of microseconds. The test stand recorded extended waveforms up to 12 µs after initial illumination, during which all pulses were recorded. Based on the method described in ref [173] the distribution of integrated probabilities for an afterpulse event up to 12 µs following an initial event is shown in Figure 3.6.

3.1.2 LED Calibration

Light from six 468nm LEDs is routed via optical fibres to various locations around the tank. These are used to measure in-situ PMT response, cable propagation delays, and perform relative transparency monitoring.

In-situ gain measurements were performed using the same method as with the test stand and found results in good agreement; the mean difference between datasets was 10% and the 1σ spread in differences was 27%. PMT voltages were tuned based on these data to a gain of 7x10^6, achieving good uniformity (within a 7% spread) and SPE separation (peak-to-valley ratios of ∼1.5–2).

The LEDs are controlled by an HVSys driver system[174], which monitors the emitted light with a PIN photodiode and uses a feedback circuit to stabilise the output. This provides illumination levels constant to within 0.5%, suitable for simple relative transparency monitoring. By tracking the number of photoelectrons generated on PMTs at the top of the tank in response to

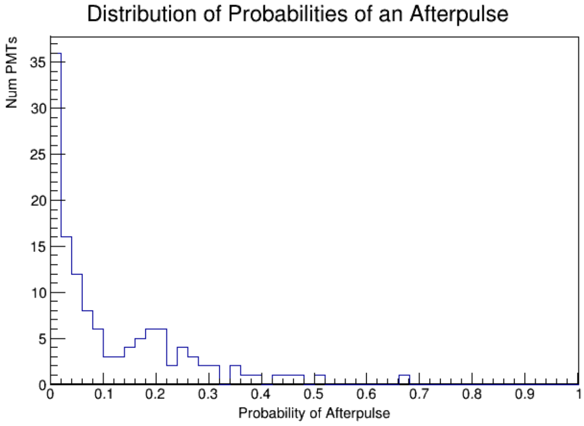

Figure 3.6: Distribution of the afterpulse probability for tested PMTs. The probability of observing an afterpulse from each initial pulse is normalized to the initial pulse amplitude, resulting in probabilities per initial photo-electron.

LED flashes at the bottom of the tank any reduction in transparency of the water can be identified. The measurement uncertainty is dominated by eventwise fluctuations in pulse charge, but changes in transparency of $\sim \mathcal{O}(2\%)$ should be detectable. This technique has been in use continuously since the water fill; readings taken over the first few months of running are shown in Figure 3.8.

The LEDs were also used to perform an initial measurement of the signal propagation time for each PMT. The physical positions of all PMTs, LEDs and calibration ports, were measured to <0.9 mm precision in a laser scan performed by Exact Metrology[175]. A snapshot of this scan is shown in Figure 3.9. Given this knowledge, the time between emission and detection of flashes from each of the six LEDs can be used to derive the signal propagation time for each PMT. The results of this measurement are shown in Figure 3.10. The measurement uncertainty is limited by the minimum 5 ns duration of the LED pulse. While this isn't precise enough for calibrating the positions of the LAPPDs, it's sufficient for a first pass for the PMTs. More accurate measurements are planned for later in commissioning and will use a PILAS laser capable of pulses as short as 45 picoseconds to provide timing calibration with the accuracy required to take full advantage of the LAPPDs.

3.1.3 Neutron Source

An in-situ measurement of neutron capture and detection efficiency crucial for calibrating the neutron yield measurement. This will be performed with an americium-beryllium (AmBe)

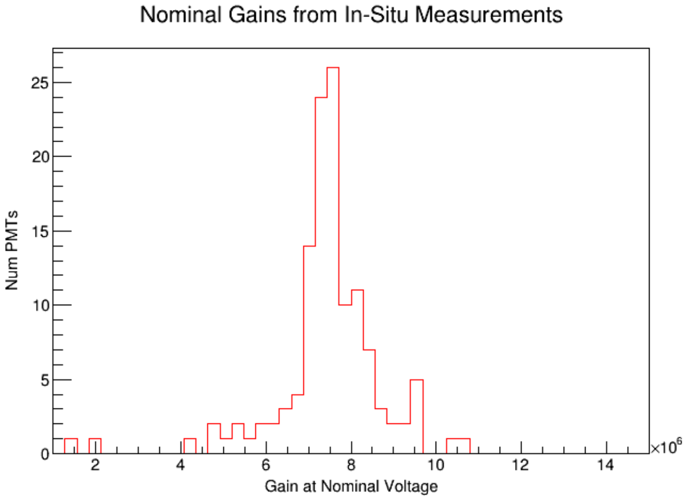

Figure 3.7: Distribution of the gains for all tank PMTs

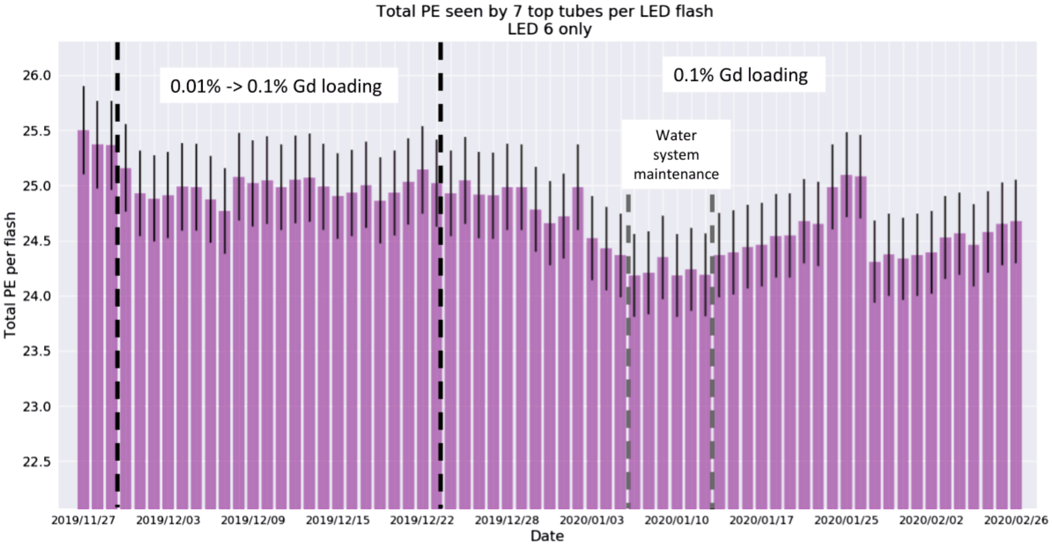

Figure 3.8: Relative water transparency throughout the gadolinium loading, estimated by monitoring the mean number of photoelectrons generated by PMTs at the top of the tank in response to LED flashes at the bottom of the tank. Note that 0.1% gadolinium represents the full loading (0.2% $Gd_2(SO_4)_3$).

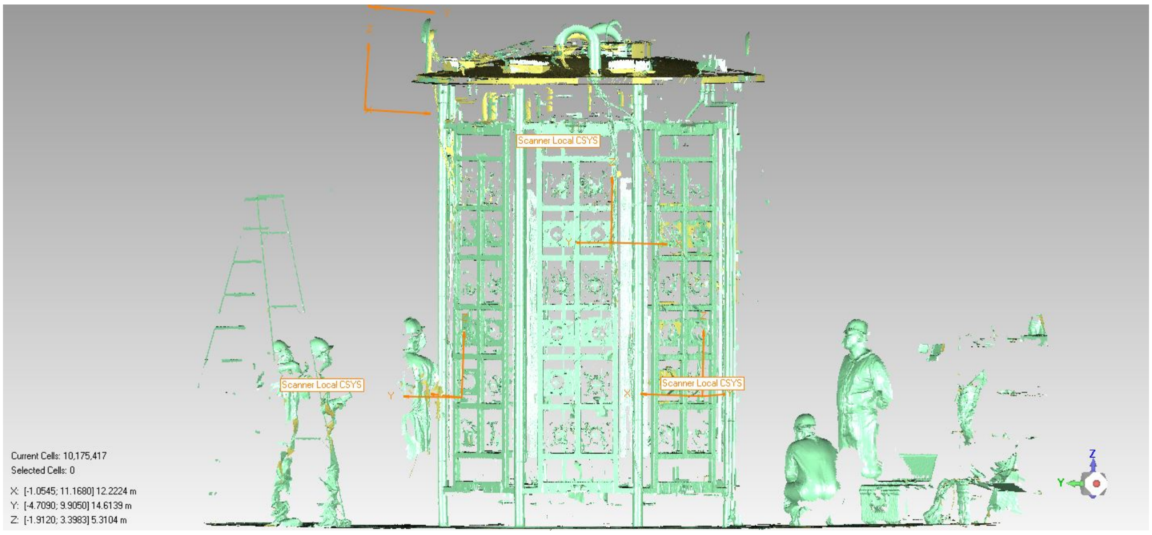

Figure 3.9: An infrared laser scan performed by Exact Metrology determined the relative position of all objects in the tank (PMTs, LEDs, calibration ports etc) to within 0.9mm. The scan used a Surphaser 410 hemispherical scanner.

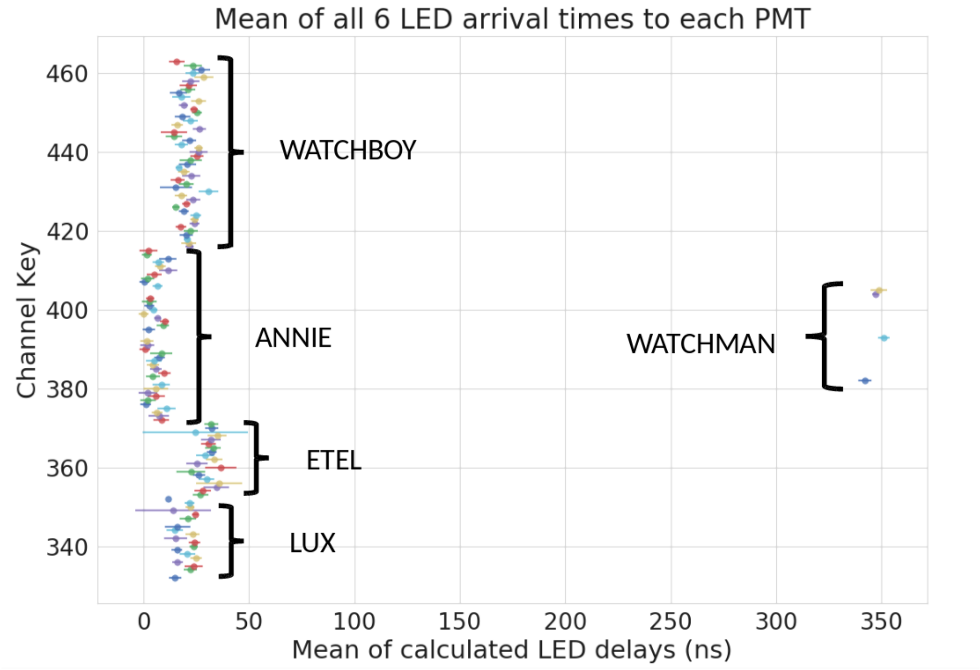

Figure 3.10: Initial measurements of signal propagation delays for the PMTs, using the LED flasher system.

neutron source coupled to a bismuth germanium oxide (BGO) crystal to provide a trigger signal on each decay. Decays of americium proceed as:

1. $^{241}\mathrm{Am} \rightarrow {}^{237}\mathrm{Np} + \alpha$
2. $\alpha + {}^{9}\mathrm{Be} \rightarrow {}^{12}\mathrm{C}^{*} + n$
3. $^{12}\mathrm{C}^{*} \rightarrow {}^{12}\mathrm{C} + \gamma$

The americium source activity is 100 µCi and has a 432 year half-life. This provides a sufficient rate of neutrons for good statistics, while being low enough to minimize pileup. The long half-life means the source rate will remain stable well beyond the experimental requirements.

The 4.4 MeV gamma emitted by the excited carbon will Compton scatter in the BGO crystal, producing an optical flash used for triggering. The BGO crystal has a light yield of 8500 photons per MeV, peaking in the 375-650 nm range, and is monitored by two silicon photodiodes. Triggering on coincidence achieves a signal to noise ratio of 25:1, an efficiency of 10^{-4} and a rate of $\sim$100 tagged neutrons per second.

The emitted neutrons have an average energy of 4.2 MeV and a maximum energy of 11 MeV, spanning a similar range as those from neutrino interactions; they should therefore have similar characteristic lengths for thermalisation and diffusion.

The source setup is encapsulated in a water-tight acrylic vessel supported by stainless steel wires to minimize shadowing. A schematic of the vessel is shown in Figure 3.11. This will be deployed at various depths through five calibration ports located on the top hatch of the tank (visible in Figure 2.10), providing multiple points with which to calibrate the detection efficiency map. Preliminary measurements have shown detection efficiencies ranging from 30–60%, based on data with an estimated 99% purity of neutron capture events[176].

3.1.4 Standard Candle

As well as the LEDs a radioactive standard candle will provide fixed energy events for calibrating PMT responses and monitoring transparency. The standard candle design, shown in Figure 3.12, consists of a 1 µCi ^{137}Cs source and two BGO crystals, and will provide consistent energy events without the need to calibrate any electronics. UVT acrylic spacers allow the separation between the caesium source and BGO crystals to be changed, varying the rate of events between 10 Hz and 25 kHz. This will also be deployed at various central locations in the tank through the calibration ports.

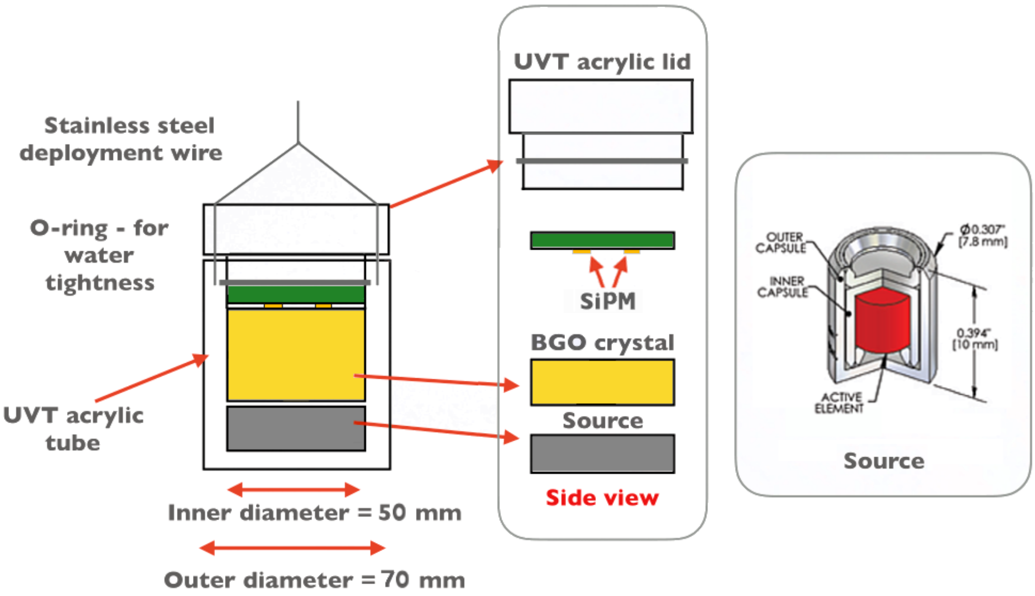

Figure 3.11: Design of the AmBe neutron source

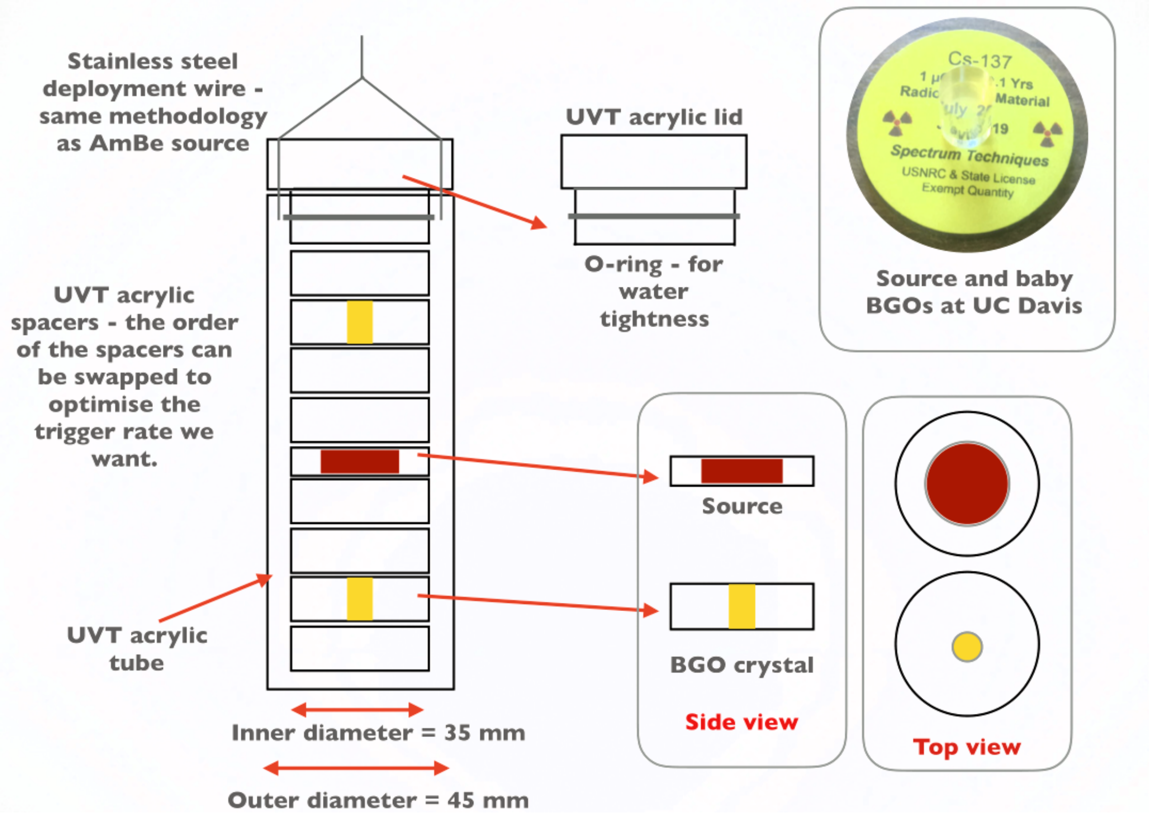

Figure 3.12: Design of the Cs standard candle source

3.2 LAPPDs

As one of the key components of the experiment, and an early adopted new technology, each LAPPD is undergoing extensive characterization before installation into the detector. This is being performed on a test stand at Fermilab shown in Figure 3.13. The setup includes a computer-controlled translation stage; a PILAS pulsed laser for timing and gain measurements; and a NIST-calibrated LED and picoammeter for QE measurements[177]. For each LAPPD the following tests are performed:

- A complete 4 x 4 x 4[1] voltage scan, recording single PE pulses while varying the voltages across each internal stage of the tile over three different values;
- The gain is measured over an 8x8 grid of positions, including a high-statistics sample at one reference point for reduced uncertainties;
- The dark rate is measured by counting the number of pulses in a large sample of random triggers with no light source;
- The quantum efficiency is measured in a 81 x 81 point grid over the entire active area of the LAPPD, once on initial power-up and again at one week without power interruption;

For a subset of the LAPPDs additional measurements of spatial resolution are taken by recording high statistics samples of 1–5 PE waveforms at regular intervals along the length of the tile, both parallel and perpendicular to the anode striplines. Some exemplary results are shown in Figure 3.14. A publication of the complete characterization procedure and findings is in preparation.

As an early adopter, ANNIE's first batch of LAPPDs includes several tiles produced while manufacturing processes were still being finalised; this primarily resulted in lower quantum efficiencies than later production units. The average QE values of LAPPDs tested by the ANNIE collaboration at the time of writing are shown in Table 3.1. Development is on-going with the tiles on hand, but it is expected that the lower QE tiles will be exchanged before installation into the tank.

---

[1] The voltage between the lower MCP bottom and anode is kept fixed at 200V, giving 4 remaining free potential differences

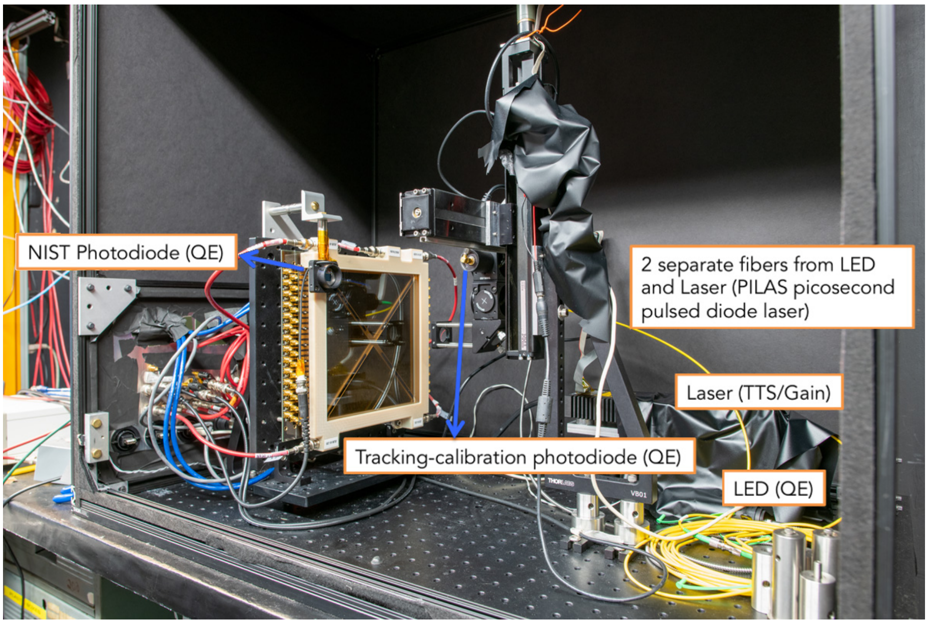

Figure 3.13: The LAPPD characterization test stand at Fermilab

ID	Mean QE [%]	1σ [%]
58	24.4	2.6
39	23.7	1.3
40	16.1	1.3
25	8.2	1.3

Table 3.1: Quantum Efficiencies of ANNIE LAPPDs, measured at 420nm.

3.3 MRD

As part of the refurbishment discussed in section 2.3 the efficiency of all new paddles installed into the MRD were tested using three-fold coincidence measurements at 3 points along their length. The results found efficiencies in the range 80–90% along the entire paddle length, for both KTeV and SciBooNE paddles. These results have since been found to be in good agreement with in-situ measurements that cover the entire MRD. In-situ measurements are performed by reconstructing MRD tracks and searching for events in which the reconstructed track passes through a paddle but no hit was recorded. A typical event, with a missing hit highlighted, is shown in Figure 3.15, and the resulting map of efficiencies is shown in Figure 3.16. The majority of paddles show efficiencies in excess of 80%; techniques to improve the efficiency of the remaining paddles are under investigation.

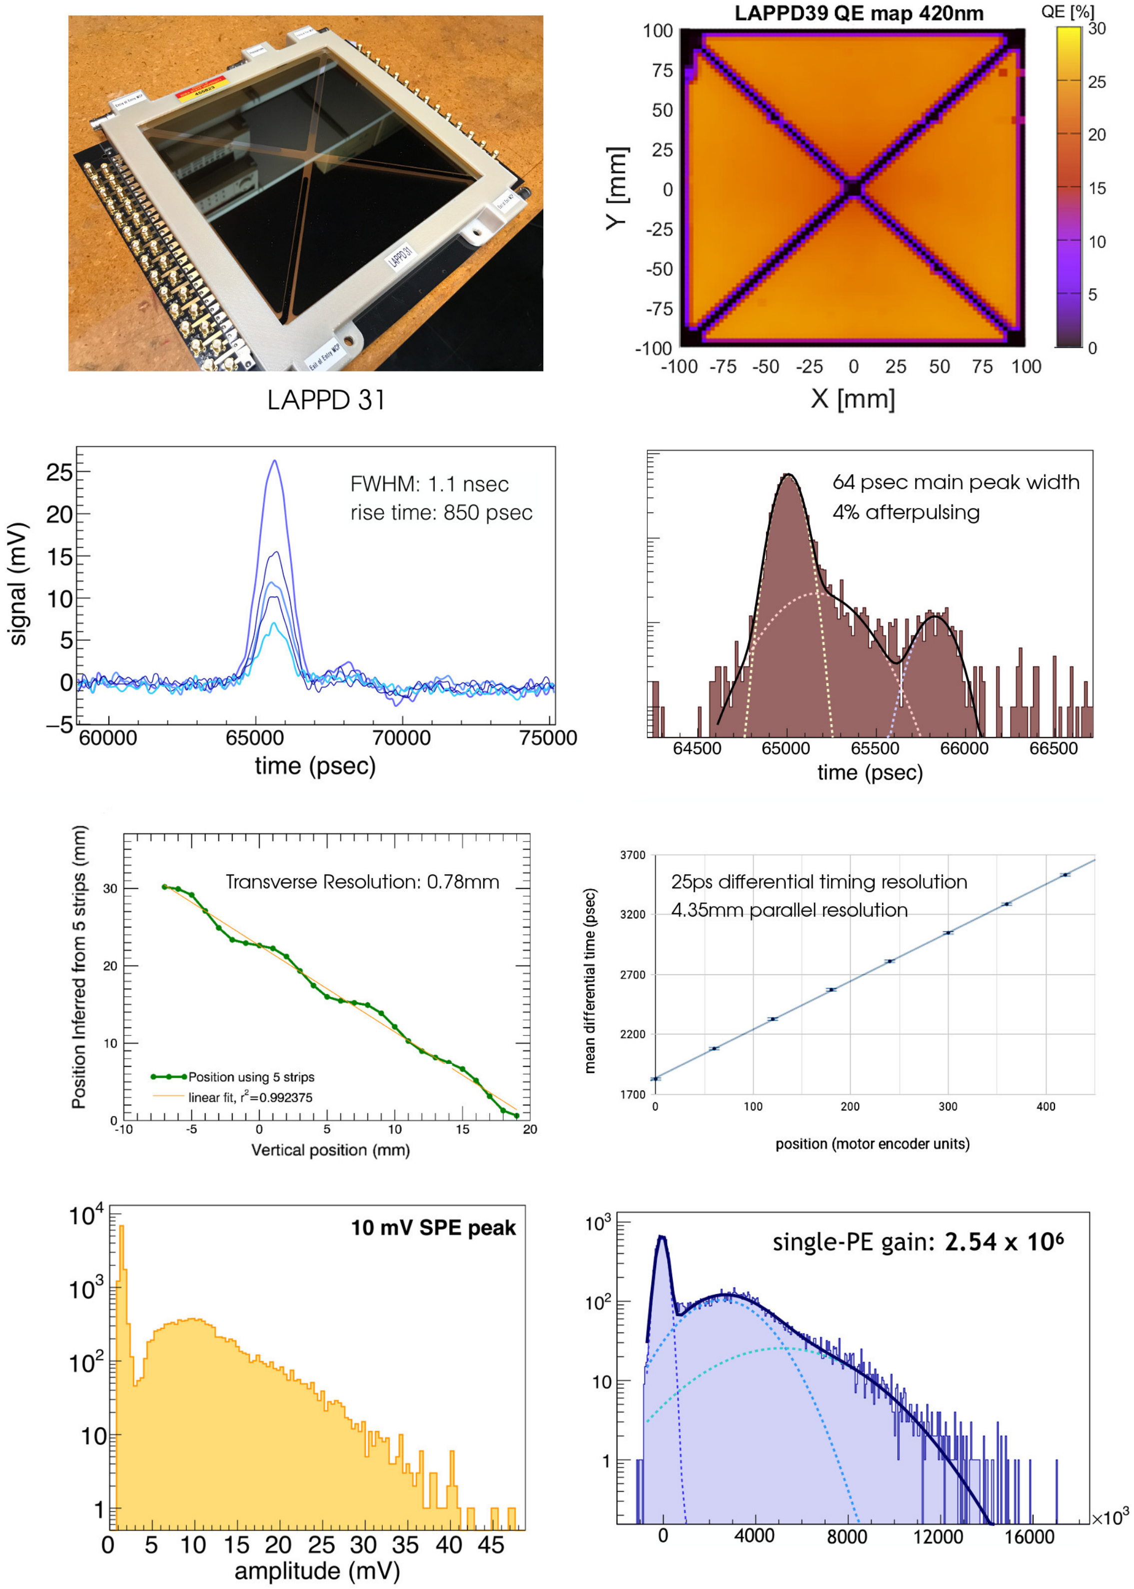

Figure 3.14: 1) an LAPPD. 2) quantum efficiency map. 3) typical pulses. 4) transit time spread for 1PE pulses. 5) charge centroid vs position transverse to anode striplines. 6) differential timing vs position along striplines. 7) 1PE pulse height distribution. 8) 1PE pulse charge distribution.

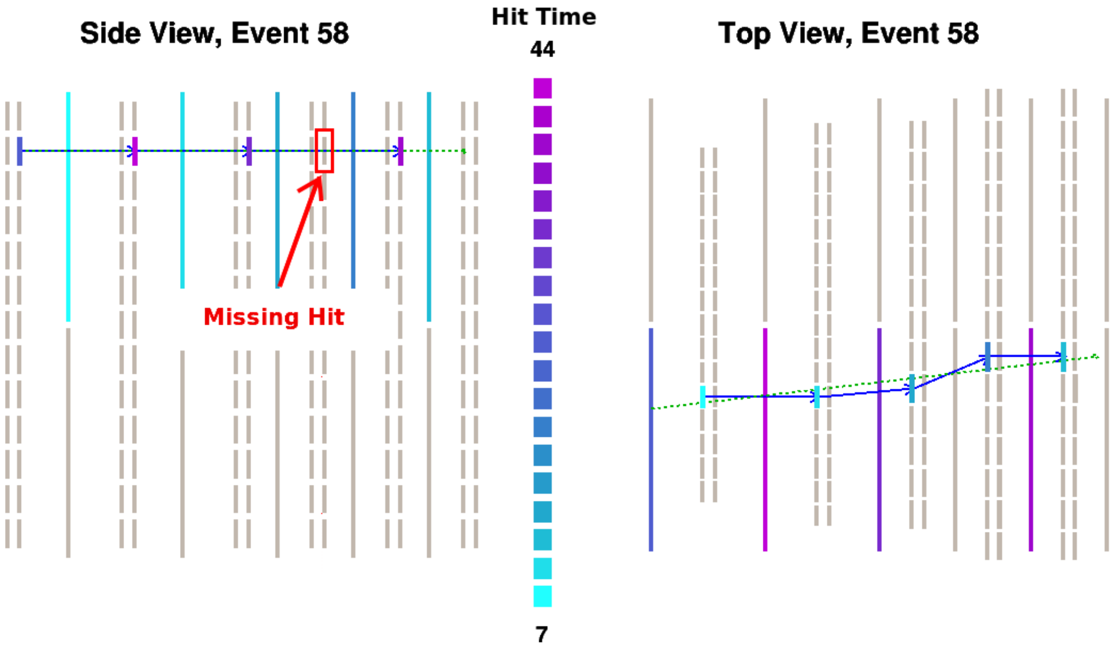

Figure 3.15: A schematic view of the MRD, with coloured paddles indicating those with a hit. The green dashed line represents the reconstructed track, with a missing hit highlighted in the 4th horizontal layer.

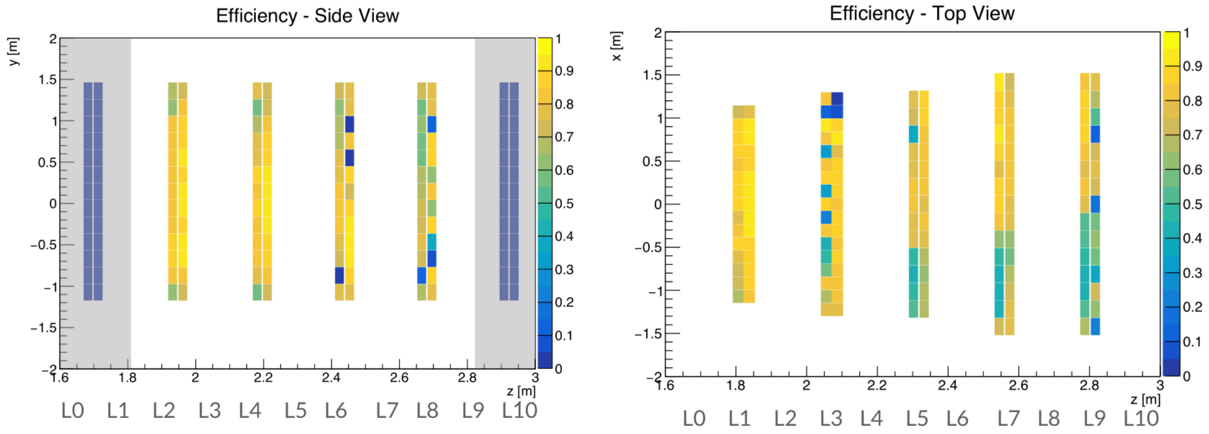

Figure 3.16: Measured efficiency of MRD paddles, based on missing hits such as those shown in Figure 3.15. Efficiencies of the first and last horizontal layers have yet to be determined as they will require the method used to be extended. Work is ongoing to address the low efficiency paddles.

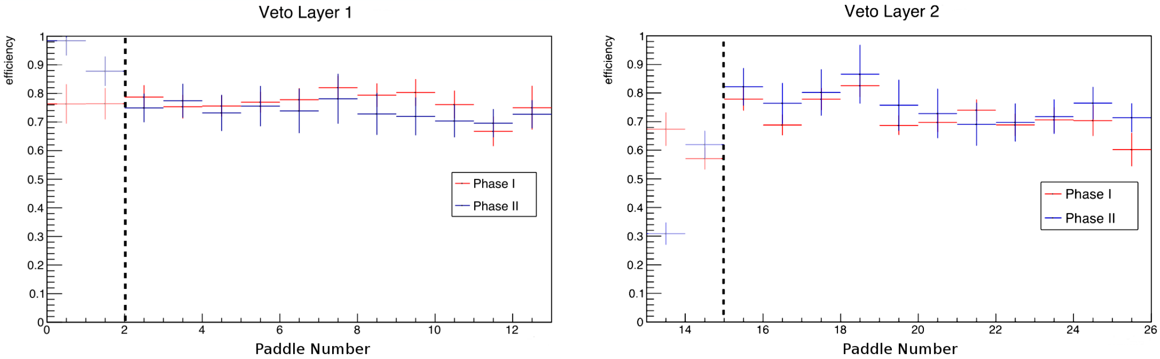

Figure 3.17: Efficiency of the first (left) and second (right) forward muon veto layers, averaged across the length of each paddle. The first two channels, to the left of the dotted lines, currently have light leaks that are expected to be fixed soon, so the efficiency for these channels is currently undetermined. Plots by Michael Neislony.

3.4 FMV

The efficiency of the forward veto is assessed using a similar method as for the MRD. Events with a track in the MRD and tank, coincident with a hit in at least one layer of the veto, are checked for a corresponding hit in the other layer. By back-projecting the track a position-dependant efficiency can be calculated across the veto wall. The efficiency of hit detection of the paddles falls linearly with distance from the PMT by up to 30%, but placing the two layers in opposing orientations and requiring an 'OR' of the layers helps to compensate for this reduction, resulting in a minimum combined efficiency of ~90% across the entire FMV. The average efficiencies of the FMV paddles are shown in Figure 3.17, using results from Phase I (red) and more recent Phase II (blue) data.

4

The Fermilab Booster Neutrino Beam

4.1 Proton Beam Generation

Neutrino flux is provided by the Boost Neutrino Beam (BNB). The beam starts life as a stream of hydrogen ions (H^-) accelerated by a 500 ft linear accelerator up to energies of 400 MeV. These ions are injected into the booster proton synchrotron ring, where a thin carbon foil strips off the electrons producing a proton beam. The booster ring is 150 m in diameter and consists of 24 repeated periods, each containing a series of 4 combined function bending and focusing magnets. Periods operate on the strong focusing principle, alternating horizontally focusing and vertically defocusing ('D') magnets with horizontally defocusing and vertically focusing ('F') magnets. When used in combination this produces an overall focusing in both vertical and horizontal axes. Each period contains two quadrupole magnets of each type, interleaved with short ('O') and long ('OO') drift sections, arranged in a 'FOFDOOD' pattern.

Seventeen RF cavities provide bunching and acceleration of the beam over a period of 33 ms. The RF cavities are excited with the 84th harmonic of the revolution frequency, producing 84 stable phase space areas known as buckets, each of which may contain one proton bunch. Of the 84 total only 81 buckets are filled, leaving a gap for the extraction kicker magnets to ramp to full voltage. Shifting the RF phase relative to the beam generates the desired acceleration. One beam revolution takes 2.22 μs on injection, falling to 1.6 μs by extraction. As the beam accelerates the bending field strength rises from 740 Gauss to 7,000 Gauss, while the RF frequency is ramped from 37.86 MHz to 52.81 MHz.

Extraction is performed by 4 kicker magnets and a septum magnet, which together provide 45 mrad of displacement that knock the beam out of orbit. A switch magnet then diverts the

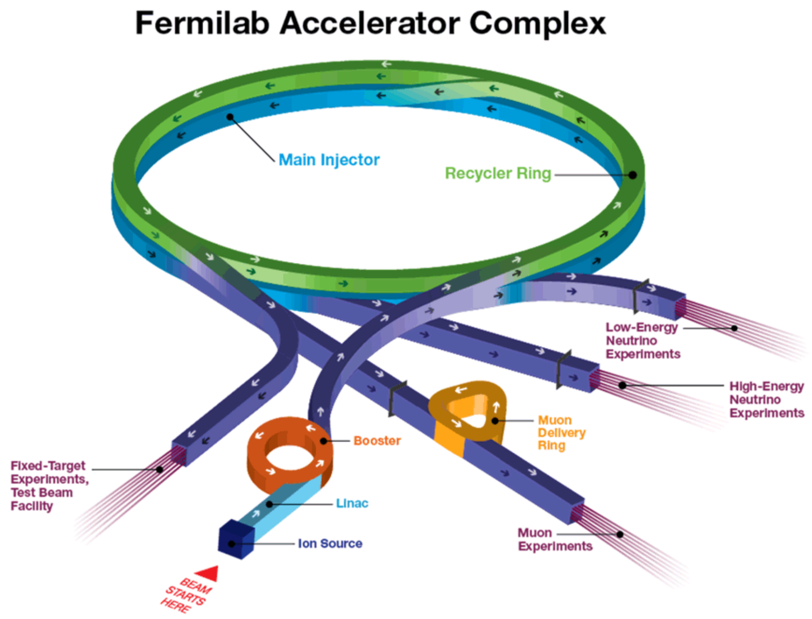

Figure 4.1: The Fermilab Accelerator Complex. The complex provides many facilities with particle beams of several different types and energies. The BNB is the beam for the 'Low-Energy Neutrino Experiments' indicated on the diagram. Figure from [180]

beam to BNB[1], followed by 3 quadrupoles that capture and focus the beam through a 42 m drift tube. A further 3 quadrupoles then match the beam to a final set of 'FODO' magnets that produce the remaining deflection toward the BNB target. Further information can be found in refs [178] and [179].

The resulting spill structure spans 1.6 µs, contains 81 bunches around 1.5 ns in length, with each bunch separated by 19 ns. Beam spills are delivered at an average rate of 7.5 Hz and each contain $\sim$4x10^{12} protons.

4.2 Neutrino Beam Generation

The 8 GeV protons extracted from the booster are directed onto a 71 cm beryllium target. Several factors make beryllium a suitable choice as the target material. The isotopes produced by proton bombardment (^{7}Be and Tritium) are relatively safe, reducing residual radioactivity concerns. Its thermal and mechanical properties allow beryllium to withstand the thermal shock and material fatigue induced by proton bombardment over a long lifespan. Finally it has a high pion production yield, leading to high neutrino yield, and a low Z that helps reduce energy lost to undesired interaction channels. The target is divded into 7 segments to allow room

---

[1] The same extraction point is also used for the Main Injector, recycler and beam dump

for thermal expansion, increase the efficiency of forced air cooling, and mean that mechanical failure of one segment does not compromise the entire target.

Proton collisions with the beryllium target generate a spray of secondary particles, consisting mainly of nucleons, pions and kaons, ejected with a wide spread of angles and momenta. The mesons, predominantly pions with energies up to 6 GeV, decay to produce the neutrinos of interest. The majority of the neutrino flux is ν_μ arising from the chain $p + Be \rightarrow \pi^+ \rightarrow \mu^+ + \nu_\mu$. This channel dominates up to neutrino energies of $\sim$ 2.8 GeV, above which the chain $p + Be \rightarrow K^+ \rightarrow \pi^0 + \mu^+ + \nu_\mu$ takes over. In total ν_μ make up 93% of the neutrino flux, with the remainder consisting of 6.4% $\overline{\nu_\mu}$, mainly from $p + Be \rightarrow \pi^- \rightarrow \mu^- + \overline{\nu_\mu}$, and 0.6% ν_e from decays of the μ^+ produced in the main π^+ channel and via $p + Be \rightarrow K^+ \rightarrow \pi^0 + e^+ + \nu_e$ at energies > 1 GeV[181].

Neutrinos are emitted isotropically in the rest frame of the parent meson, but relativistic beaming transforms this into a strongly forward emission profile. Focusing the parent mesons before decay therefore allows the formation of a neutrino beam. This focusing is performed by a toroidal aluminium 'horn' surrounding the target (top left in Figure 4.2). The horn is pulsed with 170 kA of current over a 150µs window, peaking at the time of beam arrival. This generates a toroidal magnetic field reaching 1.5 T at its centre, tuned to focus 0.5–1.0 GeV pions into parallel paths down the decay pipe. The horn polarity may be chosen to focus positive pions to produce a neutrino beam ('forward horn current'), or negative pions for an antineutrino beam ('reverse horn current'). In each case opposite sign secondaries are also defocused, helping to reduce wrong-sign contamination. The BNB is expected to run in neutrino mode for the duration of ANNIE data taking, so all values of neutrino flux given in this dissertation are for forward horn current. More information on the target and horn can be found in [182].

The focused pions pass through a collimator into a 50 m long decay pipe ending in a steel and concrete beam stop. The decay pipe length is chosen to allow the majority of pions to decay, while being short enough that most muons from these decays are absorbed by the beam dump. This reduces the $\nu_e + \overline{\nu_\mu}$ beam contamination from muon decays.

The main components of the neutrino generation process are shown in Figure 4.2. Between the end of the beam pipe and the ANNIE detector hall is 50 m of rock. The locations of two buildings, together with those of other experiments on the Fermilab neutrino campus, are shown in Figure 4.3.

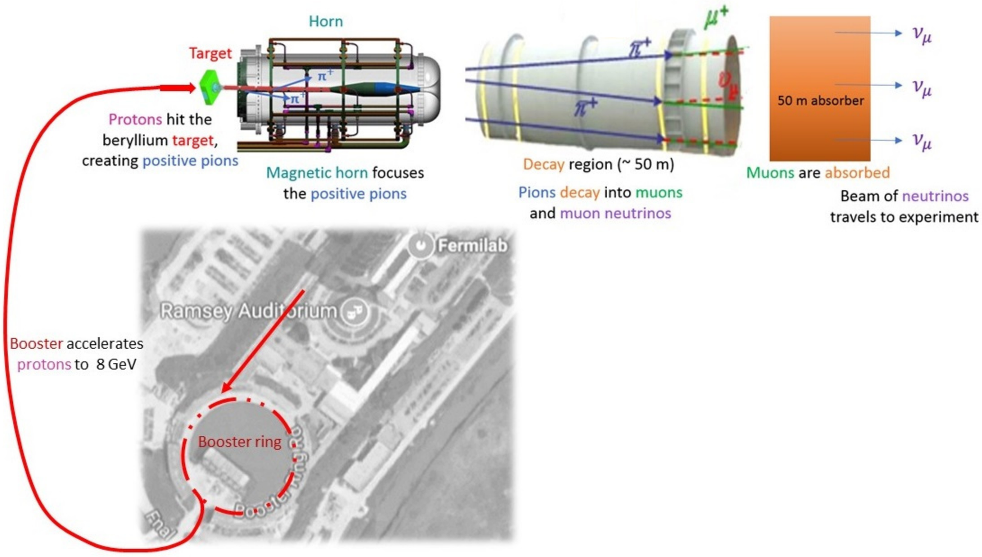

Figure 4.2: Overview of Booster Neutrino Beam generation. Figure from [183]

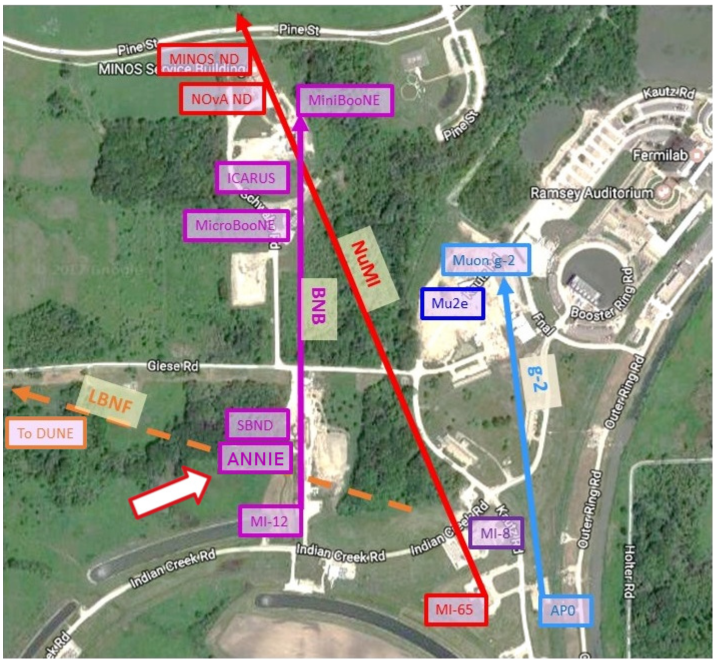

Figure 4.3: Overview of the neutrino campus at Fermilab. The BNB target hall is marked as MI-12; ANNIE is the first experiment to receive beam, immediately adjacent to the Short Baseline Near Detector (SBND). Figure from [184]

4.3 Beam Monitoring

The incident proton beam is continuously monitored by two beam position monitors, a beam profile monitor, two toroidal current monitors and a resistive wall monitor. The beam position monitors measure the charge induced on a split plate as the beam passes, and their combined output allows determination of the beam position to within 0.1 mm. A wire chamber measures the beam profile with 48 horizontal and vertical wires at a 0.5 mm pitch. Together these measurements indicate that 99.8% of protons in the beam are expected to intercept the target. Two toroidal current monitors determine the intensity of the beam by measuring the magnetic field induced in a ferrite ring around the beam pipe. The toroids agree to within 2%, and routine calibrations give results consistent to within 0.5%. This information is recorded in a database that is later used by experiments for precisely tracking the delivered flux. Finally the resistive wall monitor (RWM) is a device consisting of a ceramic insulating section of the beampipe, across which several resistors are placed. As the beam propagates it induces a current in the beam pipe, which flows through these resistors as the RWM is passed. By monitoring the voltage across these resisters the beam passing can be recorded. The RWM provides a timing signal that is forwarded to experiments, and can be used to match neutrino events to specific buckets within spills.

5

Monte Carlo Simulations

5.1 Hadron Production

5.1.1 Scattering Cross-Sections

The Booster Neutrino Beam has been extensively modelled by the BooNE collaboration[181], and this model continues to be used for neutrino flux predictions by other experiments on the BNB line. Likewise ANNIE will be using these simulations to predict the neutrino flux at the experimental hall.
Predictions begin with a simulated flux of protons 4.5 cm upstream of the BNB target, with initial position and angle randomly selected from a span matching the beam optics. Simulations are handled by the Geant4 framework, but the cross-sections for primary proton and secondary pion interactions with the target and horn, together with the selection of products from these interactions, are overridden with a custom model based on fits to experimental hadron production data. An analysis of the discrepancies between Geant3, Geant4 and the MARS15 packages motivating this replacement can be found in [185]. In the MiniBooNE model interaction cross-sections are broken up into three main components; coherent elastic scattering off the nucleus as a whole, quasi-elastic scattering off a single nucleon in which the particle is not absorbed and no other hadrons are produced, and the remaining 'reaction' cross-section that encapsulates all other interactions.

Total cross-sections for nuclear interactions are derived from those for free nucleon interactions, for which considerable data is available. To do this a Glauber model parameterisation[186] is used to first obtain the amplitude for forward elastic scattering. The forward elastic region is used as it is where the model assumptions are most valid.

The Glauber model parameterises these amplitudes with the form:

$$f(q) = \frac{(i+\alpha)k\sigma_n}{4\pi} e^{\beta t/2} \tag{5.1}$$

where q is the momentum transfer, $t = |q^2|$, k is the wavenumber of the incident hadron, σ_n is the total cross-section for free nucleon scattering, and α is the ratio of real and imaginary parts of the forward-scattering amplitude. Values for these parameters are obtained from the literature. The nucleon amplitudes are then summed coherently over suitable radial distributions (Saxon-Woods for the horn aluminium, independent harmonic oscillator for target beryllium) to obtain the total amplitude for the nucleus as a whole. Finally this is related via the optical theorem to the total cross-section for the nucleus[187]. For pion scattering at low energies, where Δ resonance dominates, the total cross-section prediction transitions to using a Breit-Wigner parameterisation of data in this region.

Total inelastic cross-sections are also well-measured across the full range of proton momenta and the majority of relevant pion momenta. An empirical parameterisation is fit to these data, inspired by Regge theory[188] but with parameters that have no strict physical meaning.

The inelastic cross-section is further split into quasi-elastic and reactionary channels. Correctly accounting for the quasi-elastic component is important as the outgoing nucleon maintains the majority of its momentum and can undergo further scattering. Around 7% of muon neutrinos from π^+ decays originate from protons that have undergone multiple scattering events.

Little data exists to perform this separation, so the shadowing model is used to extend free nucleon quasi-elastic cross-sections to those for nuclei, accounting for attenuation of the wavefunction while traversing the nucleus. The same parameterisation is used for the separate quasi-elastic and reactionary components. More information can be found in ref [181].

Elastic scattering cross-sections are taken as the difference between the Glauber model predictions of the total cross-section and the sum of inelastic cross-sections. The final breakdown is shown for protons in Figure 5.1 and for pions in Figure 5.2.

5.1.2 Secondary Production

Measurements of the production cross-section for π^+, π^-, K^+ and K^0 are also available over the majority of the kinematic phase space relevant for the BNB. Data for the production of π^- and K^0 are fit with a Sanford-Wang parameterisation[190], which is extrapolated when necessary to cover the full kinematic phase space. In the case of the π^+ data a Sanford-Wang parameterisation resulted in a poor fit to the available data, so a spline fit was used instead[191]. For K^+ production, where data is only available at higher energies, Feynman scaling is used to extrapolate the data down to the energy range of interest[192].

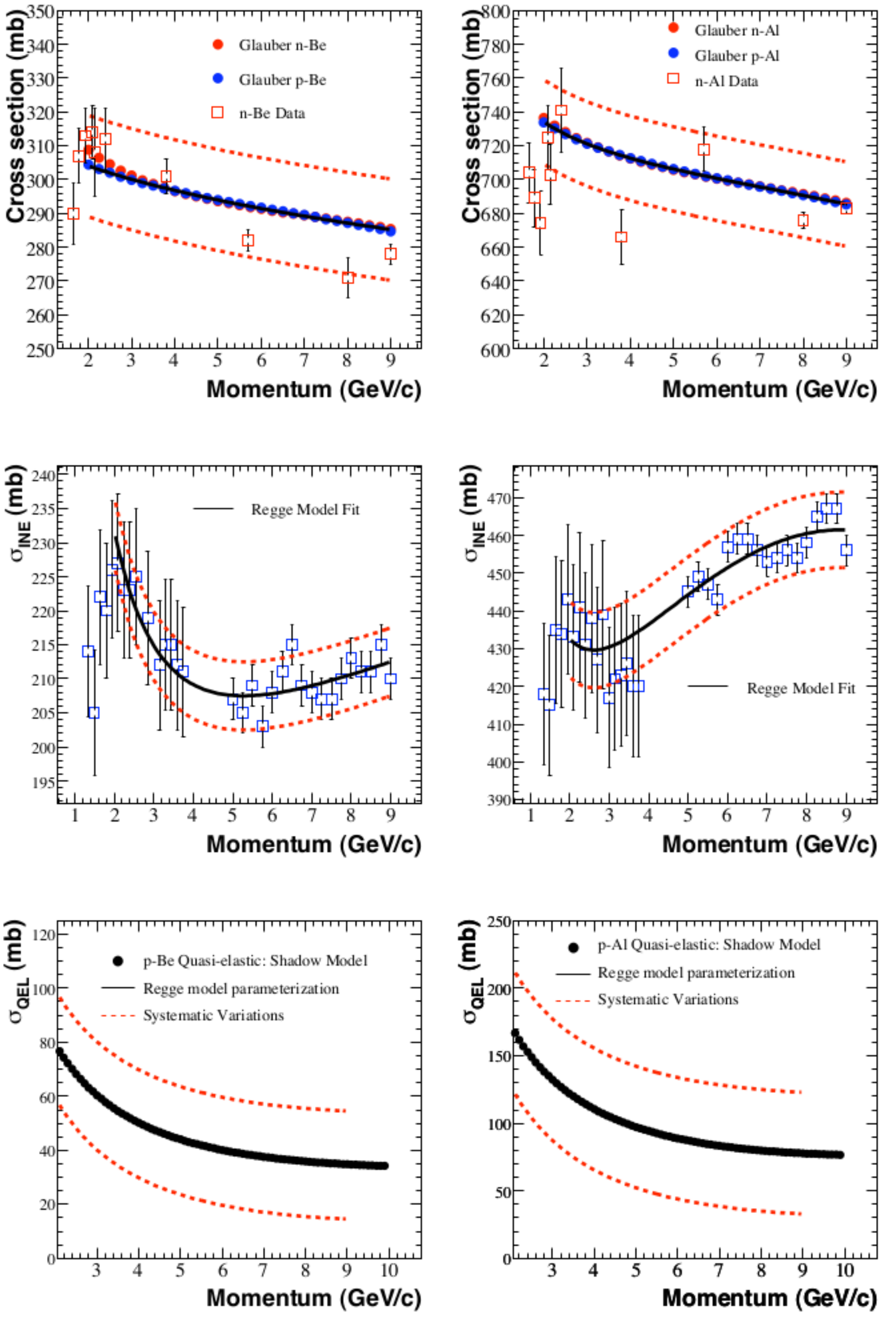

Figure 5.1: Total, reactionary, and quasi-elastic cross-sections for nucleons on beryllium and aluminium. Figure from [189].

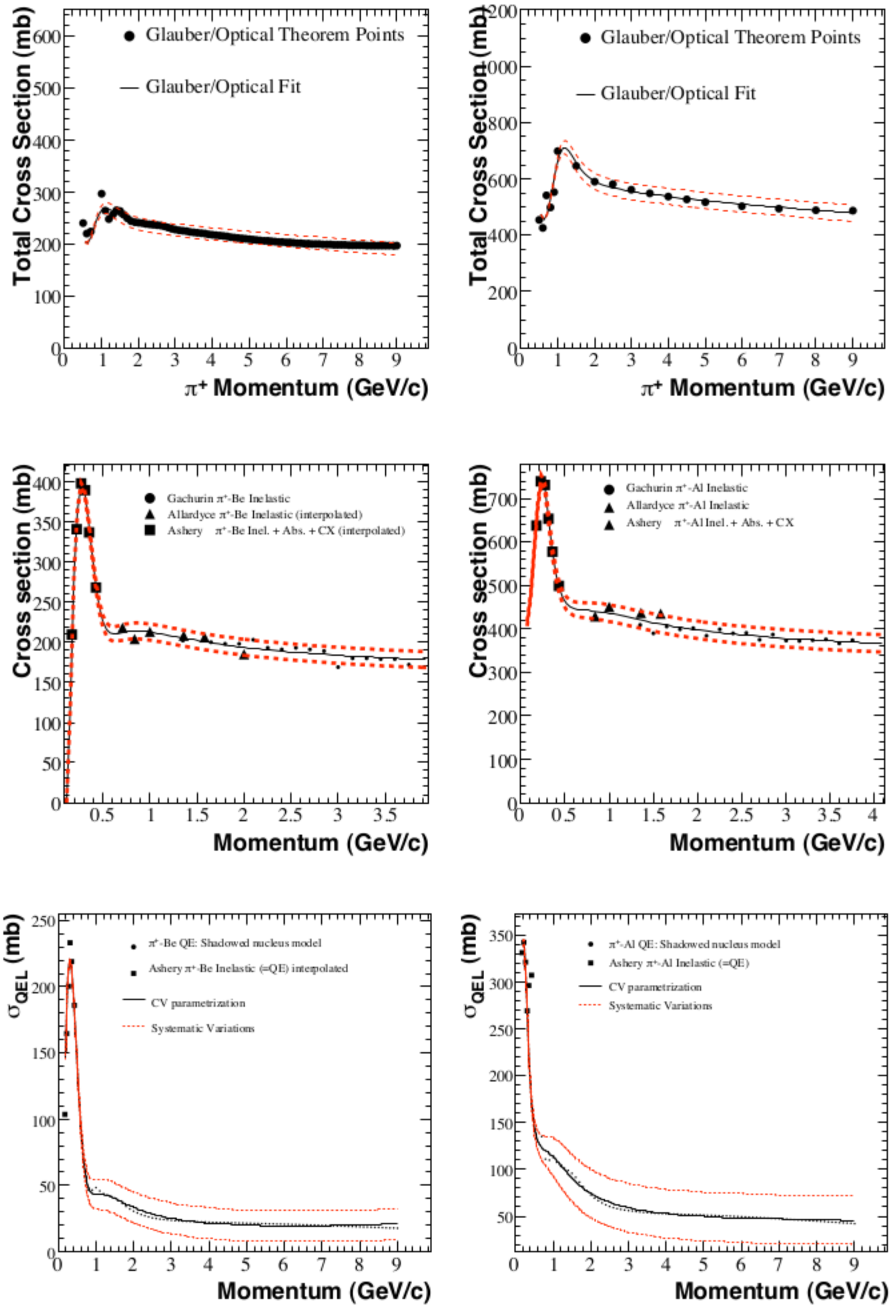

Figure 5.2: Total, reactionary, and quasi-elastic cross-sections for pions on beryllium and aluminium. Figure from [189].

Insufficient data is available for such a fit to production cross-sections for secondary protons, neutrons and K^-, so they are instead based on the MARS14 hadronic interaction package. All multiplicity distributions are double differential as a function of incident proton momentum in longitudinal and transverse directions. The number of particles of each type produced in each interaction is obtained by sampling Poisson distributions with averages based on the ratio of production cross-section for that particle type to the total inelastic cross-section. Secondary kinematics are obtained from normalised distribution functions in solid angle and momentum. Energy and momentum are not necessarily conserved on an event-wise basis, but the resulting average fluxes and spectra should be correct. A summary of the resulting particle multiplicities and average kinematics are given in Table 5.1.

Particle	**Multiplicity**	$\langle p \rangle$ [GeV/c]	$\langle \theta \rangle$ [mrad]
p	1.5462	2.64	441
n	1.3434	1.59	586
π^-	0.9004	0.82	556
π^+	0.8825	1.11	412
K^+	0.0689	1.69	332
K^0	0.0241	1.34	414
K^-	0.0024	1.26	259
Total	4.7679	1.69	496

Table 5.1: Particle multiplicities and average kinematics for each non-quasi-elastic primary proton interaction. Values from [181]

5.1.3 Neutrino Production

The branching ratio of particle decays were based on 2006 data from the Particle Data Group[119]. The values used are shown in Table 5.2. Only those neutrinos that point towards the detector are recorded. This would result in a high computational expense for the desired statistics, as many events would produce no output. To improve the efficiency each hadronic decay is simulated 1000 times, each time resampling the daughter kinematics. The neutrino flux is weighted against the number of protons simulated to account for this enhancement. Other secondary channels with a particularly low probability are further enhanced. Each muon from secondary decays is simulated 20 times, since most will be stopped by the beam dump before they have a chance to decay. This improves the statistical uncertainties on contamination from those muons that do decay. For the high energy tail where statistics are also low this kind of re-sampling is not appropriate. This is because the parents are highly forward boosted, resulting in a narrow spread of daughter kinematics. Re-sampling the daughter kinematics would tend to obtain similar results each time, but the resulting sample would still be granular due to the low numbers of parents. To ameliorate this high energy meson production is enhanced by scaling the production cross-section up by a factor exponential in longitudinal momentum. This artificially enhances the production of high-energy mesons, with the corresponding neutrinos then weighted down to account for this. More information can be found in ref [189].

Particle	**Lifetime [ns]**	**Decay Mode**	**Branching Ratio [%]**
π^+	26.03	$\mu^+ + \nu_\mu$	99.9877
		$e^+ + \nu_e$	0.0123
K$^+$	12.39	$\mu^+ + \nu_\mu$	63.44
		$\pi^0 + e^+ + \nu_e$	4.98
		$\pi^0 + \mu^+ + \nu_\mu$	3.32
K^0_L	51.6	$\pi^- + e^+ + \nu_e$	20.333
		$\pi^+ + e^- + \overline{\nu}_e$	20.197
		$\pi^- + \mu^+ + \nu_\mu$	13.551
		$\pi^+ + \mu^- + \overline{\nu}_\mu$	13.469
μ^+	2197.03	$e^+ + \nu_e + \overline{\nu}_\mu$	100.0

Table 5.2: Decay modes of the secondaries that contribute to the neutrino flux. Decay characteristics of respective negative particles are the same so they are not shown. Data from [181]

5.2 Neutrino Flux at the ANNIE Detector

All neutrinos that intercept a 20x20 m window 20 m upstream of the experimental hall are recorded for the next stage of simulation. This window covers the hall together with a region of the surrounding dirt. Interactions in the dirt are used to estimate the flux of secondary particles from external interactions that enter the detector. Of particular importance are interactions that produce neutrons, as these constitute a beam-correlated background for the neutron yield measurement. This background is discussed in detail in subsection 5.4.1.

The predicted neutrino flux spectrum is shown in Figure 5.3. The total neutrino flux is estimated to be 2.2 x10^{-8} cm^{-2} per proton on target. The flux is dominated by muon neutrinos, amounting to 93% of the total neutrino flux, while muon antineutrinos contribute 6.4% and electron neutrinos just 0.6%. The contributions from each decay channel to the muon neutrino flux is shown in Figure 5.4.

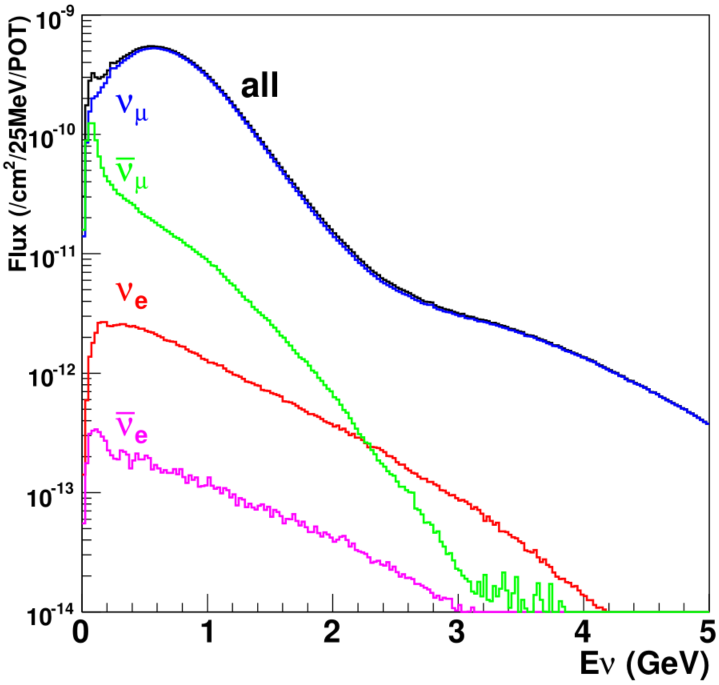

Figure 5.3: The neutrino flux at the ANNIE experimental hall. Figure from [193].

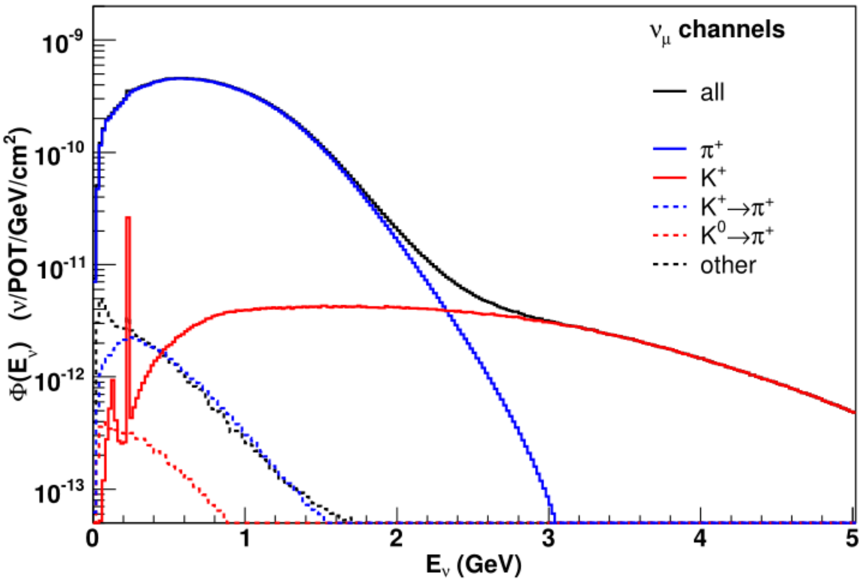

Figure 5.4: The fractional contribution to the muon neutrino flux from each contributing decay mode. Figure from [181].

Systematic errors on the flux prediction are divided into two categories. Event-wise systematics, such as uncertainties in the production cross-section of a parent meson, can be accounted for by event re-weighting. Every neutrino has a weighting factor for each systematic associated with its generation. The nominal flux is obtained by assuming the central value for all weights. To see how a given systematic uncertainty affects the flux it is sufficient to re-sum the events with the corresponding weight contribution varied appropriately.

Other uncertainties are not specific to the event, such as variations in the focusing horn current. To obtain the variation in flux from these uncertainties 1000 different 'universes' are simulated, with parameter values spanning the full range of the uncertainty. In each case the flux spectrum is calculated, and the final flux uncertainty is taken from the variation over all spectra.

The largest uncertainties arise from particle production cross-sections, particularly that of π^+ as the dominant source of muon neutrinos. Averaged over all neutrino energies the resulting uncertainty in muon neutrino flux is $\sim 14\%$, likely to be a dominant factor in cross-section measurements. These uncertainties are significantly inflated to account for the inability of the parameterisation to match the data, suggesting better models are needed more than higher precision measurements.

A comparison of the results from these models with MiniBooNE data found an observed rate of ν_μ interactions $\sim$20% higher than nominal predictions. While this discrepancy is covered by the prediction uncertainties, data driven normalizations would be valuable in reducing the uncertainties on any cross-section measurements.

5.3 Neutrino Interactions

Neutrino interactions are simulated with the GENIE v2.12.0 generator with the `Default` event generator list configuration[194]. The initial nuclear state is modelled by a Bodek-Richie Fermi gas - a global relativistic Fermi gas model with an ad hoc extension above the Fermi momentum to account for short range correlations[195][196].

Quasielastic scattering is modelled with the Llewellyn-Smith formalism, with the BBBA05 parameterisation of nucleon form factors. This parameterisation writes the weak vector form factors in terms of the electric and magnetic form factors of the proton and neutron, which are then extracted from a piecewise fit of polynomials to electron scattering data[197].

Resonant pion production is calculated with the Rein-Seghal model, extended to account for lepton masses. Only the 16 lowest mass resonances are included and the model does not take into account the interferences between resonances. Coherent pion production is also calculated with the coherent Rein-Seghal model.

	NC	**CC**	**CCQE**	**CC-Other**
All	11323	26239	13764	12565
With Muon Entering MRD	2	7466	4279	3187
With Muon Stopping in MRD	2	4830	2792	2038
With Muon Fully Penetrating MRD	0	1454	761	693
With Muon Exiting the MRD Side	0	1181	726	455

Table 5.3: Number of events per year within a nominal 2.5-tonne fiducial volume of the tank

Deep inelastic scattering, including all non-resonant pion production, is handled using the Bodek-Yang model. Hadronisation of interactions with invariant masses below 2.3 GeV are modelled by a custom AKGY model[198], transitioning to Pythia[199] above 3 GeV.

Multinucleon (np-nh) interactions are not included.

Final state interactions are handled by an effective cascade model, which estimates the final state particles for each initial state particle in one step. The model is tuned to data of proton and pion scattering on iron, and extrapolated to other nuclei based on $A^{2/3}$ scaling. The advantage offered by this technique is that nuclear effects are easily reweightable.

These models represent the default configuration offered by GENIE 2.12.0. Many alternative models are available within the framework but as yet no tuning has been performed. The chosen models may change as data becomes available and analyses progress.

Neutrino interactions are simulated within a simplified GDML geometry that includes the detector, experimental hall and surrounding rock. Roughly 26,000 charged-current interactions are expected per year in a 2.5-tonne fiducial volume of the water tank. Around 5,000 of these events are expected to produce a muon that stops in the MRD, allowing complete reconstruction of the event. A breakdown of the expected event rates is given in Table 5.3. The number of neutral current interactions is of a similar order to those of charged current interactions. These intteractions can result in the ejection of neutrons, representing a potential background source of neutron production[200]. The low rate of pileup is crucial to suppress this background, with the low rate of interactions giving a probability of both charged and neutral current neutrino interactions in the same beam spill of $< 10^{-6}$.

5.4 Background Sources

5.4.1 Dirt Interactions

The products of simulated neutrino events are initially propagated by a simple Geant4 model that transports the particles until they either range out, leave the simulated world extent or intercept the ANNIE detector. Only those products that intercept (or originate within) the detector are recorded for the next step of simulation. As well as making the detector simulation more efficient this step also provides a measurement of the flux of particles from external interactions.

While the majority of muons from dirt interactions should be vetoed by the upstream scintillator wall, any that do slip through represent a background for CC0π or CC-inclusive cross-section measurements. Simulations indicated that 0.7% of beam spills produced a muon from an external neutrino interaction that intercepted the water tank, with around 80% of those passing through the veto.

Unlike muons, neutrons can be particularly problematic as they produce no signal in the forward veto and represent an irreducible, beam-correlated background for the neutron yield measurement. The distribution of neutrino vertices producing a neutron that enters the tank is shown in Figure 5.5. Around 1% of beam spills produce a neutron entering the tank, 79% of which arrive within the 70µs recorded event window. As the chance of capturing a stray neutron is independent of the chance of a neutrino interaction, this means $\sim$0.79% of events with a neutrino interaction are expected to also contain a dirt neutron within the observation window.

5.4.2 Neutron Background Measurement

Relying on the yield of neutrons from simulated events to estimate the background of a neutron yield measurement would be fairly circular. More importantly, the dominant source of background neutrons is actually not modelled by these simulations - so called ‘skyshine neutrons’. This phenomenon represents a flux of neutrons ejected in primary interactions in the BNB target, which leave the target hall at a high angle and scatter off the atmosphere before raining back down[201][202]. This phenomenon was previously observed at the site of the ANNIE detector by the SciBooNE collaboration. They reported a rise in the rate of hits at the top of the detector over the course of the beam spill, as shown in Figure 5.6. These hits are believed to come from skyshine neutrons.

To evaluate the rate of background neutrons, from both skyshine and dirt interactions, an initial phase of the ANNIE experiment ran from June 2016 to September 2017. A simplified

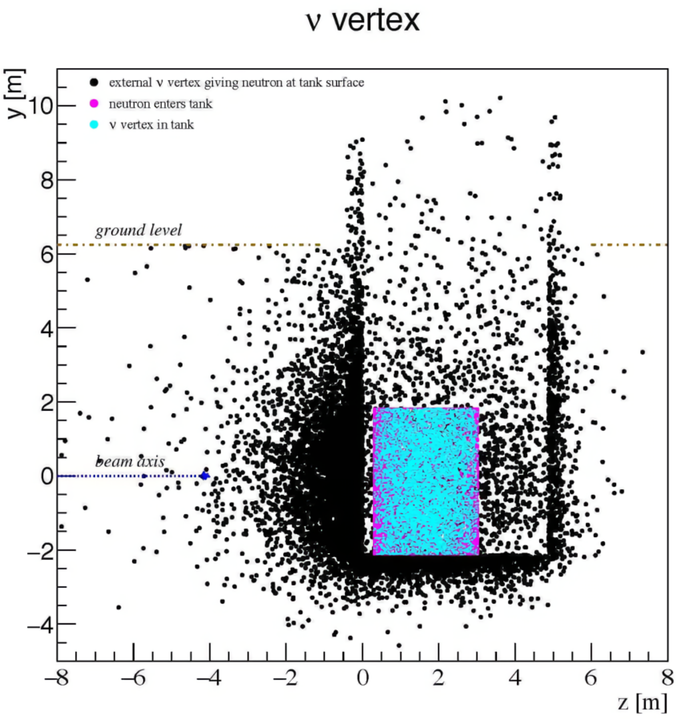

Figure 5.5: Dirt interactions producing neutrons that enter the ANNIE detector.

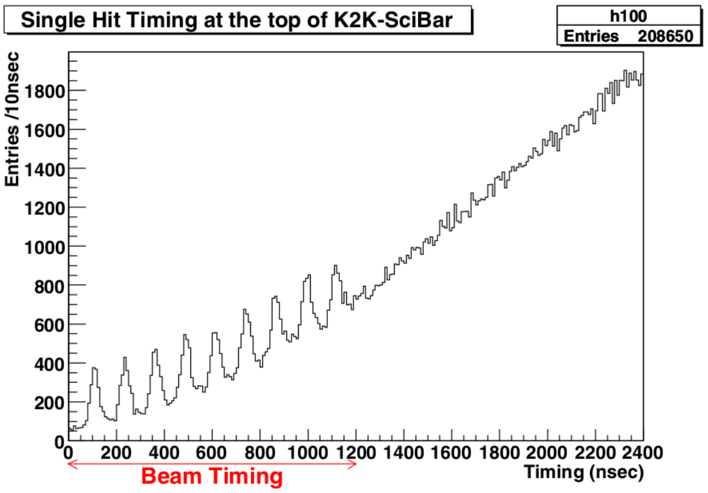

Figure 5.6: A measurement of hit rate in the top layers of the SciBar detector by the SciBooNE collaboration. The initial pulsing arises from the beam spill structure, but the steady rise is believed to come from skyshine neutrons. Figure from [203].

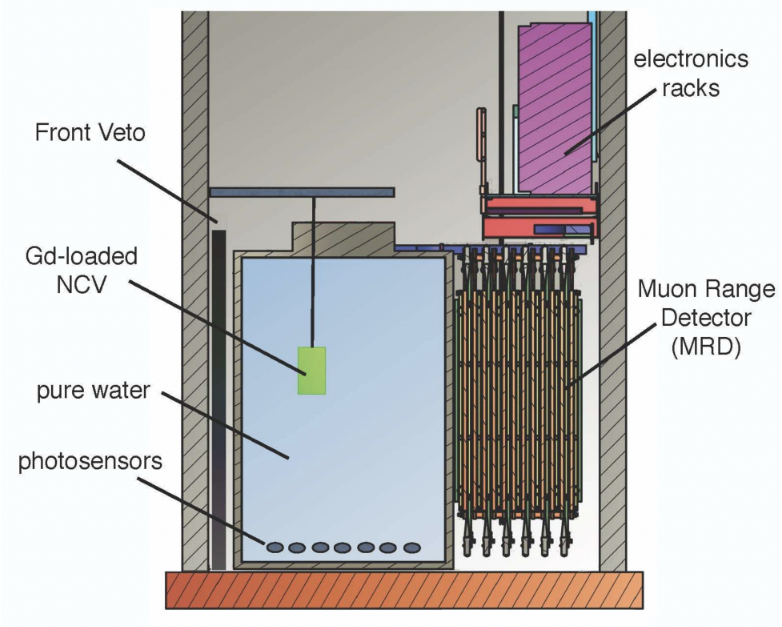

Figure 5.7: The detector configuration used in the Phase 1 measurement of background neutron flux.

detector configuration was used as illustrated in Figure 5.7. The forward veto and MRD were not used and only sixty 8" Hamamatsu R5912 PMTs were installed, all on the bottom of the tank, which was filled with pure water (i.e. no gadolinium). No blacksheeting was used, which together with the white inner liner helped to maximize the light collection efficiency. In this case scattering was not a concern as the measurement did not require interaction type classification or kinematic reconstruction - the tank PMTs were simply used as a veto. Instead the volume of interest was a 24" long, 20" diameter acrylic vessel filled with $\sim$ 100 litres of gadolinium doped liquid scintillator. This sub-volume was optically isolated from the water tank and observed by two dedicated 3" PMTs (see Figure 5.8).

The measurement strategy was to count the number of neutron captures within this subvolume that were not associated with a Cherenkov flash from a muon in the water. The sub-volume was movable, allowing measurements to be taken at 5 different depths and 3 different offsets from the front wall of the tank. By varying the amount of water shielding for dirt and skyshine neutrons, the contribution from each source could be assessed. A summary of the results are shown in Figure 5.9. The findings indicated a significant dominance of skyshine neutrons. They also demonstrated that the flux drops to less than 1/3 of that at the surface with just 15 cm of water shielding, indicating a soft energy spectrum. The optically isolated volume used in the upcoming physics run provides 36 cm of water overburden and 20–27 cm of upstream water

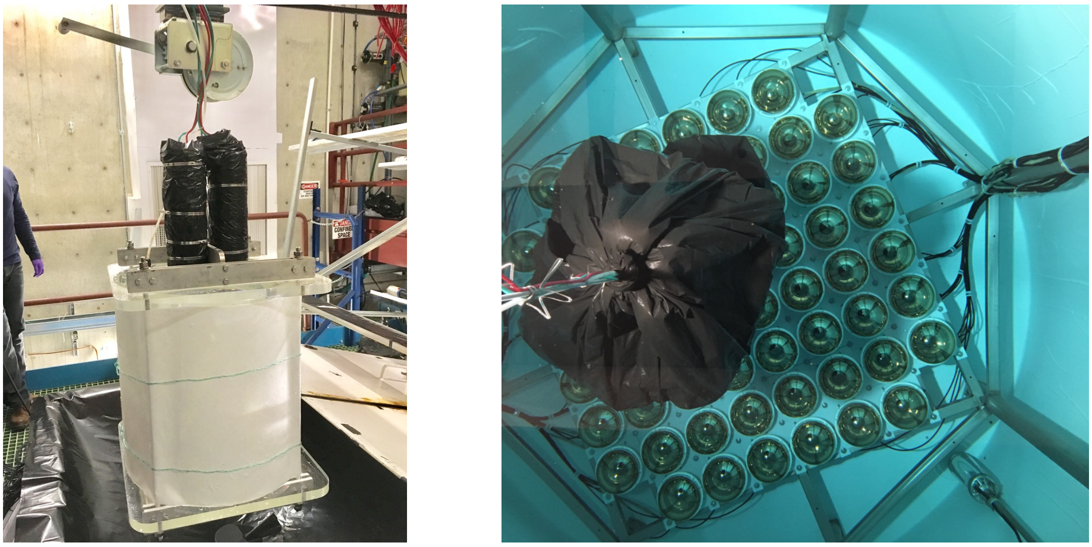

Figure 5.8: The active volume was an optically isolated acrylic vessel of EJ-335 liquid scintillator, with a high gadolinium doping of 0.25% by weight. The use of a liquid scintillator helped improve detection efficiency, and the high doping concentration helped improve capture efficiency within the smaller volume. An inner reflective wrapping improved the light collection, while an outer black wrapping minimized light leakage.

burden, providing more than adequate shielding. Averaging over the 2.5-tonne fiducial volume, just 0.02 neutrons are expected per cubic meter per spill. More information can be found in ref [115] and the thesis of Steven Gardiner[204].

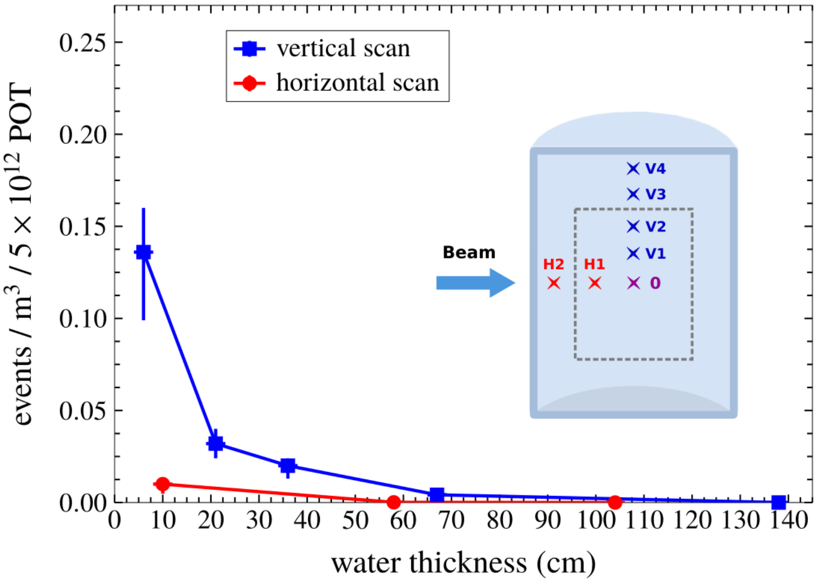

Figure 5.9: A summary of background neutron flux, measured in Phase I of the experiment.

5.5 Detector Simulation

Detector response is simulated by a Geant4[205] model built with the WCSim framework[206]. WCSim provides a starting point for developing Geant4 simulations of water Cherenkov detectors, including classes for PMT objects, hit digitization, multiplicity triggering, and the construction of a scalable water tank. From the base code offered by WCSim various extensions were required to fully implement the ANNIE detector. Primarily these focused around adding support for the additional detectors external to the water tank, along with their respective hit collections, digitization and triggering schemes. Modifications were also required to accommodate the many photosensor types used in the main water tank, their asymmetrical placement, the steel support structure, and the digitization scheme required for simulating full waveform readout, among others.

The simulation uses Geant4 version 10.2 configured with the `FTFP_BERT_HP` physics list[207]. This is the current default physics list, with the option to enable high precision modelling of low energy neutrons, including libraries specialized for scattering of thermal neutrons in water. Neutron interaction properties are taken from the G4ENDL4.5 dataset, which is derived from the data in the ENDF/B-VII.1 library[208]. This library provides tables of experimental data and standardized parameterisations for neutron interaction cross-sections, secondary production yields, energy and angular distributions. The resulting predictions have been validated against both experiment and the MCNPX simulation package[209][210].

WCSim tracks the interactions of each particle through the detector, including the production of Cherenkov and scintillation light, its absorption, reflection and scattering. The optical characteristics of surfaces in the detector are specified with the UNIFIED framework[211] to ensure reflection and absorption is properly modelled.

Photons striking a PMT bulb may reflect, be absorbed, or may pass into the PMT glass. While within the glass photons may be absorbed, resulting in hit generation. PMT quantum efficiencies are applied in two stages; an initial cut during photon generation uses the maximum quantum efficiency of any sensor in the detector, minimizing the number of photons that need to be tracked. On each hit a second cut is then applied if necessary to account for the difference in quantum efficiency of the particular sensor being hit. The charge for each photon is pulled from a random distribution based on the calibration measurements discussed in chapter 3. Digitization cuts are then simulated to affect accurate multiplicity triggering; hits with a drawn charge above the digitization threshold are added to a separate digit collection.

Triggering matches the scheme being used by the experiment; an initial prompt trigger is always recorded, after which further triggers are based on a count of digit multiplicity within a sliding window. Triggers may also be generated by the MRD for simulated cosmic ray studies. Once a trigger is generated all hits within the trigger window are recorded - whether or not they passed the digitization cut. This ensures that lower charge hits, which may not have met the online digitization threshold, are still recorded and will be present in simulated waveforms. Dark noise is also added, with rates taken appropriately for each PMT type. For each hit the true time and simulated charge is recorded, along with parentage information for truth studies. LAPPD hits also store positional information within the tile and are written out to a separate file. This later gets passed to a dedicated application for simulating stripline responses; keeping these stages separate allows rapid iteration of the LAPPD simulations without requiring regeneration of the Monte Carlo sample.

Following simulation the outputs of WCSim are converted to the same data format produced by the real detector. This process produces simulated waveforms, adding in electronics noise and a low frequency baseline oscillation observed in real data. It also introduces an approximate beam timing profile. Simulated PMT pulses are based on a templated pulse shape, with template parameter values drawn from distributions obtained by fitting a sample of real pulses.

Tuning of the simulation to match detector response is still in very early stages, with ongoing analyses comparing hit patterns from LED flashes, neutron capture time distributions, and energy deposition characteristics. As an example, one early study calculated the 'visible energy' of muons from neutrino interaction candidates and compared the results with Monte Carlo simulations of beam CC0π interactions. The visible energy is a simplified reconstruction of muon energy based on the sum of the MRD and tank energy losses. MRD energy loss is calculated using the procedure described in section 6.4, while tank energy loss is estimated using a linear conversion from the measured tank charge. Most muons from beam neutrinos are minimum ionising particles with $\beta \simeq 1$. For such particles the energy loss per unit length and the number of Cherenkov photons per unit length are both roughly constant. Given a measurement of the number of Cherenkov photons one can determine the total energy lost in the tank. To make this mapping one has only to find the constant of proportionality. This is done using a sample of through-going muons. Such events can be easily selected by requiring a single MRD track that back-projects through the centre of the tank and coincides with a forward veto hit. Such tracks have a well-defined length equal to the diameter of the tank, and so have a well defined energy loss. Measurements found a total tank charge tightly peaked around 4500 photo-electrons for such events, amounting to 12 photo-electrons per MeV. Returning to neutrino interaction candidates, applying this conversion and adding the MRD energy loss gives the total visible energy of the event. The comparison of total visible energy to true muon energy is shown in Figure 5.10. The shape and mean show reasonably good agreement, albeit with a wider spread to the data and a small high energy tail. As a first assessment agreement between simulations and data is within expectations; over the coming months a comprehensive tuning programme is anticipated to bring such discrepancies into alignment.

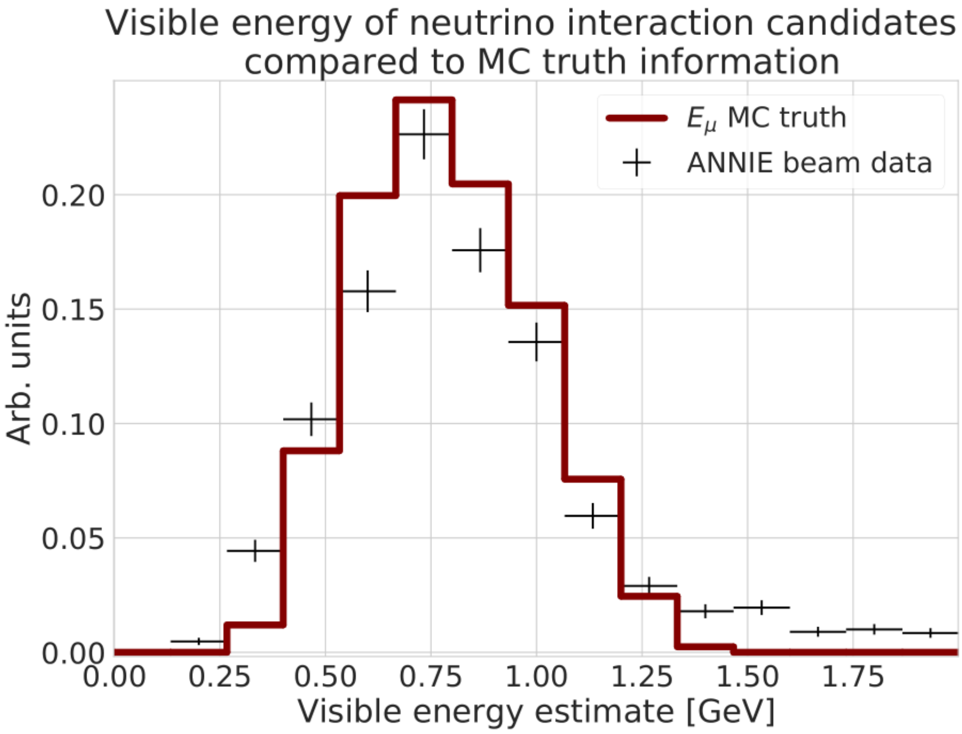

Figure 5.10: Comparison of the visible energy in events with a single MRD track and no forward veto hits with Monte Carlo simulations of beam CC0π events in the tank. The agreement helps demonstrate that beam simulations produce observable distributions that match reasonably well with real data. Figure by Teal Pershing[176].

A simulated 1 GeV muon event originating in the tank is shown in Figure 5.11. Hit PMTs are coloured according to their charge, illustrating the Cherenkov disk and MRD track. A similar event from the actual detector is shown in Figure 5.12.

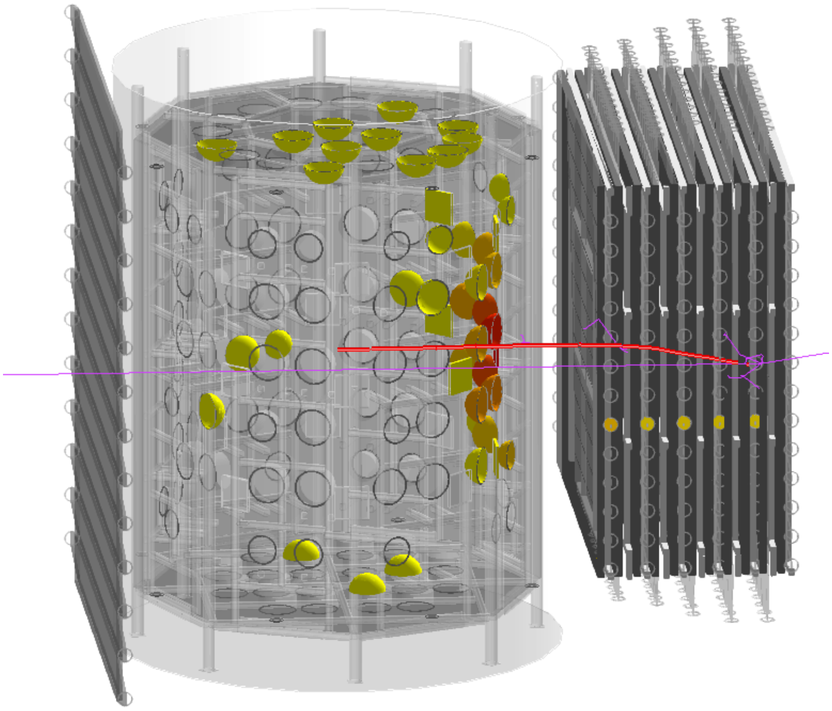

Figure 5.11: A 1 GeV muon event in WCSim. PMTs and LAPPDs with photon hits are coloured according to charge; red indicates higher charge, yellow lower charge. The red line represents the trajectory of the primary muon, while the purple lines are secondary gammas. Struck MRD paddles are also highlighted by the row of PMTs at their ends. (The highlighted PMTs appear offset from the muon track due to parallax.)

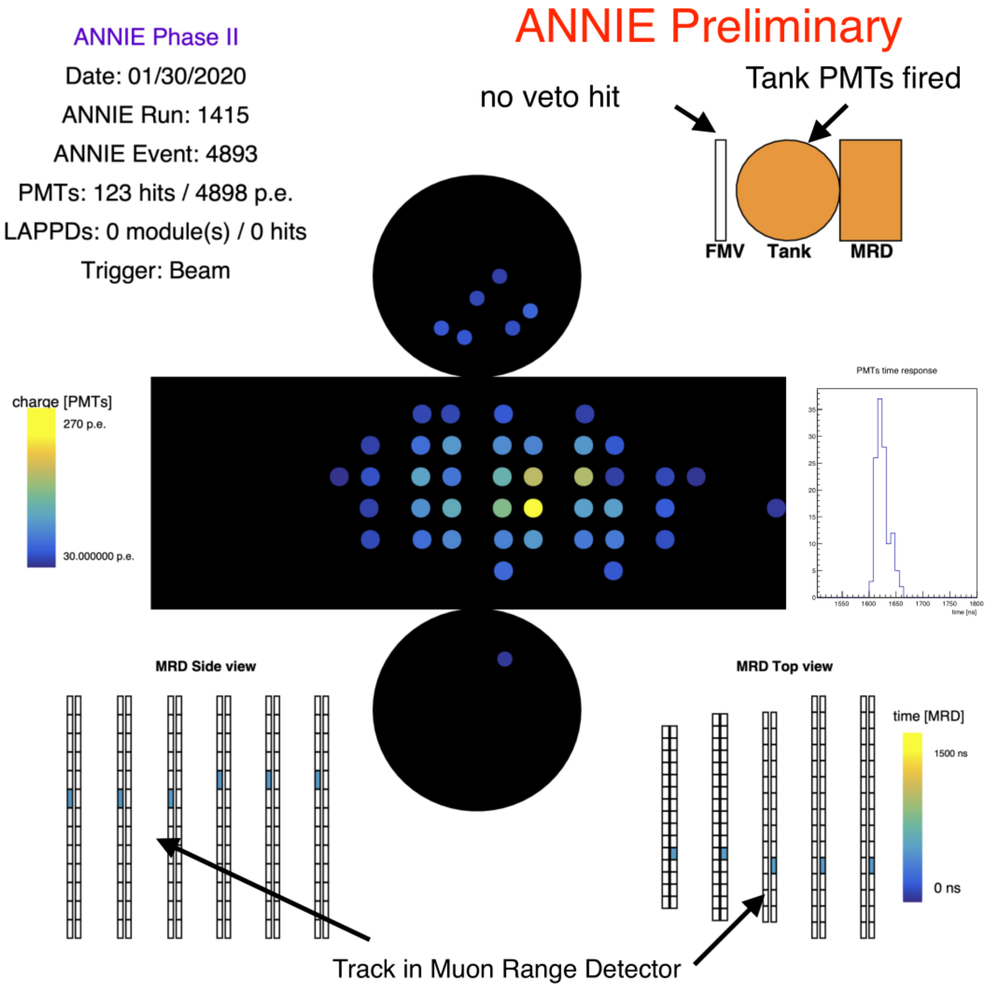

Figure 5.12: One of the first real data neutrino candidate events from the ANNIE Phase II detector. A clear Cherenkov disk is seen in the water tank, with a penetrating muon track through the MRD. The absence of hits in the upstream veto indicates this is a candidate neutrino event. A low charge cut of 30 pe has been applied to remove noise and reflections, producing a clearer Cherenkov disk.

6

Reconstruction

Reconstruction and analysis are being developed within ToolAnalysis, a fork of the ToolDAQ framework described in section 2.6 repurposed for data analysis. An initial set of tools perform event building, combining the outputs from each of the subsystems and performing timestamp decoding, alignment and data re-structuring. Waveforms are converted from ADC counts into voltages, baselines are subtracted, and PMT pulses are identified. These pulses serve as the hits for further reconstruction.

An overview of some of the most upstream tools is shown in Figure 6.1. Along with the above these tools produce a simple event display and various validation plots. This enables data quality to be assessed quickly and routinely, allowing early identification of potential problems. Note there are two toolchains shown; the upper toolchain performs data decoding and restructuring, building the `ANNIEEvent` files that will be used for offline analysis. The other loads these files (or may include the same tools as part of its toolchain) and performs the first stages of analysis.

6.1 Time Clustering

The first stage of reconstruction is a time clustering algorithm that splits the hits from the full 70µs window into distinct time-clustered events, each of which is then processed independently by the remaining tools. Time clusters are found by scanning a sliding 50 ns window over the full readout and searching for the position that encapsulates the greatest number of hits. A cluster is formed from those hits, and then the search is repeated with all remaining hits. This repeats until the largest cluster found contains fewer than 10 hits. This process separates hits that are likely to arise from a single particle, allowing each track to be reconstructed separately. As a prompt event may be followed by several delayed events, a global super-event reconstruction will re-combine events arising from a common physical interaction later.

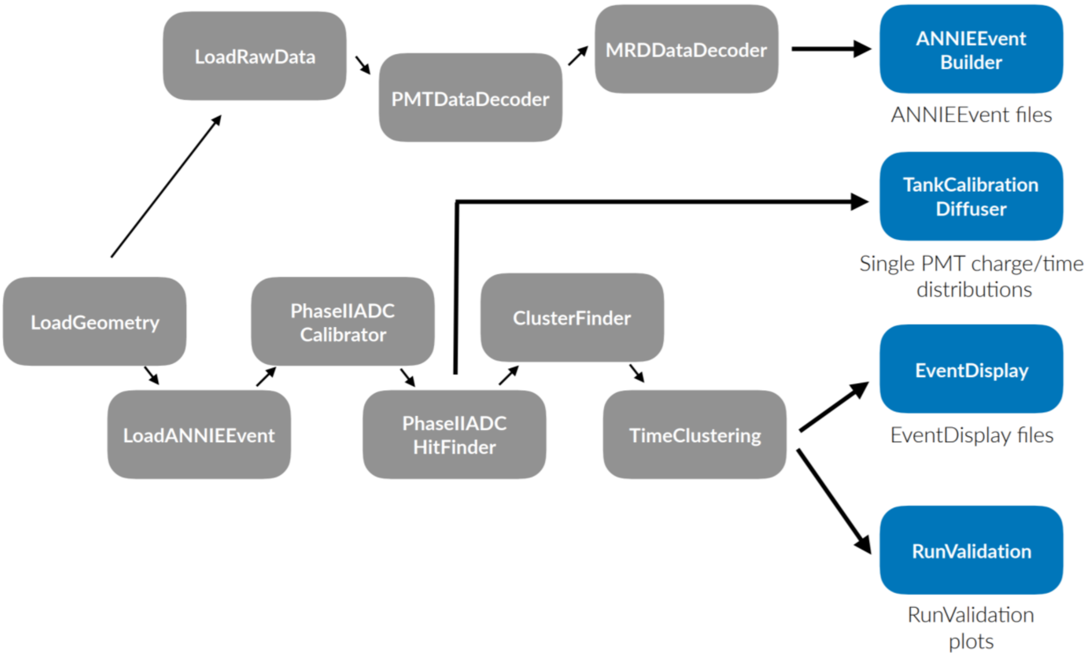

Figure 6.1: An overview of some of the upstream tools in the online monitoring and reconstruction toolchains.

6.2 Hit Cleaning

Some initial hit cleaning has proven to be effective in improving the reconstruction of vertices and directions, as well as providing a first indication of events with more than one charged particle above Cherenkov threshold. Hit cleaning begins with a simple charge cut to remove the majority of dark noise hits. A neighbouring algorithm then iterates over each hit, counting the number of nearby hits at around the same time. All PMTs within 100 cm are checked over a 10 ns window from the reference hit time. Any hits with fewer than 4 neighbours are pruned. This helps to remove higher charge hits in isolation or small groups, reducing the event to a few large clusters. Finally, the remaining hits are checked again with this same neighbouring algorithm, this time pruning any hits with no neighbours at all in a smaller 50 cm radius. The wider scan area of the first pass allows irregular cluster perimeters, while the second pass removes isolated hits near to the edge of a larger cluster. The results of each stage of cleaning can be seen in Figure 6.2. The presence of more than one cluster is a robust indication of an inelastic event.

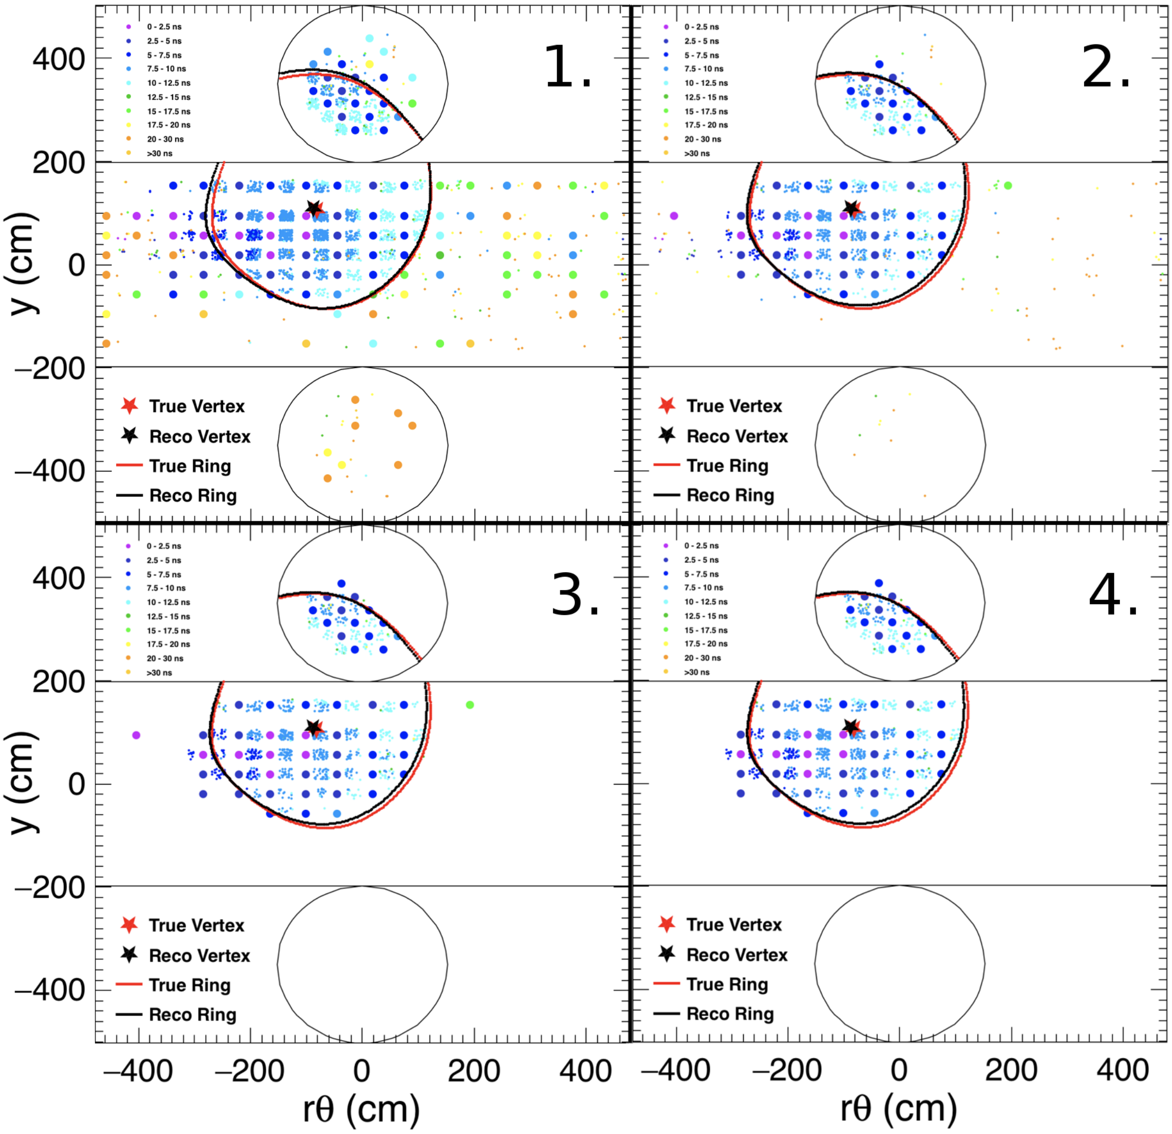

Figure 6.2: Event displays showing the stages of hit cleaning. 1) All hits. 2) After low charge cut. 3) After first neighbouring cut. 4) After second neighbouring cut. Overlaid are the true and reconstructed vertex and Cherenkov disk edge.

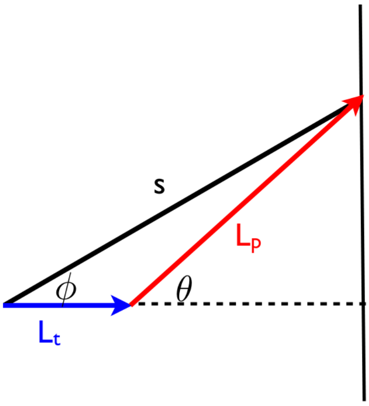

Figure 6.3: Decomposition of the hit time residual for an extended track.

6.3 Vertex Reconstruction

One of the first tools for physics analysis is vertex reconstruction. This pertains to both the interaction vertex and the direction of the ejected particle, giving six parameters, $(x, y, z, t, \theta, \phi)$, to find. Reconstruction begins with an array of seeds, regularly spaced to fill the entire fiducial volume. Each seeds time is determined from the median of all hit times after subtracting the time for propagation from the vertex to the PMT. The initial direction is calculated from a charge-weighted average of all hit positions relative to the origin. The full fitting procedure is then run over each of these seeds, and the best result is taken.

For each starting value a fully comprehensive fit to all six parameters is performed using the Minuit minimization library[212]. The quantity extremized is based on two metrics, with the first relating to hit time residuals. These time residuals represent the difference between the predicted hit time, derived from the reconstructed vertex, and the measured hit time. If the reconstructed vertex is accurate then the time residuals for all hits should coincide within the measurement uncertainties.

For a point source the expected hit time is simply its distance from the source divided by the speed of light in water. For an extended source the predicted time splits into two components; the time between particle generation and photon emission, and the time between photon emission and detection. Using the variables shown in Figure 6.3 and the speed of light in water v_c, the time residual of hit i is:

$$t_i = t_i^0 - \frac{s \cdot sin(\theta - \phi)}{c \cdot sin(\theta)} - \frac{s \cdot sin(\phi)}{v_c \cdot sin(\theta)} \tag{6.1}$$

The extent to which a vertex equalizes the residuals is quantified by the following figure of merit:

$$G = 100 - 5\chi^2 \tag{6.2}$$

where

$$\chi^2 = \frac{1}{N_{hit}} \sum_{i=0}^{N_{hit}} -2\log(P_i) \tag{6.3}$$

and

$$P_i = \frac{1 - P_n(q_i)}{\sqrt{2\pi}\sigma_i} \exp\left[-\left(\frac{t_i - \langle t\rangle}{\sqrt{2}\,\sigma_i}\right)^2\right] - P_n(q_i) \tag{6.4}$$

N_{hit} is the number of hits, σ_i is the photosensor time resolution, $\langle t\rangle$ is the mean time residual, q_i is the hit charge and $P_n(q_i)$ is the probability of observing a noise hit with charge q_i.

While the hit time residuals provide a tight constraint in the transverse direction they leave ambiguity in the interaction vertex along the direction of the track. This is because there is no penalty for extending the track further upstream - there are simply no hits for this part of the track. Some constraint may be gained along this line by fitting the cone edge. This is done by splitting the charges inside and outside the Cherenkov cone, and weighting each hit based on its distance from the cone edge. Defining $\Delta\theta$ as the angle between the hit and the projected cone edge from the start of the track, such that $\Delta\theta < 0$ for hits inside the cone and $\Delta\theta > 0$ for those outside of it, charges are weighted according to:

$$\begin{aligned}
Q_{in} &= \sum_{i=0}^{i=N_{hit}} \left(0.75 + \frac{0.25}{1 + (\Delta\theta_i/21)^2}\right) q_i, \quad \Delta\theta_i \leq 0 \\
Q_{out} &= \sum_{i=0}^{i=N_{hit}} \left(0.0 + \frac{1.0}{1 + (\Delta\theta_i/3)^2}\right) q_i, \quad \Delta\theta_i > 0 \\
Q_{all} &= \sum_{i=0}^{i=N_{hit}} q_i
\end{aligned} \tag{6.5}$$

This weighting factor is plotted as a function of angle from the track in Figure 6.4; the function peaks at 42°, corresponding to the maximum opening angle of a Cherenkov cone in water. The cone edge figure of merit is then defined as:

$$F = 100 \cdot \frac{Q_{in} + Q_{out}}{Q_{all}} \tag{6.6}$$

Extending the track further upstream would move hits toward the centre of the cone, resulting in a smaller weighting and a lower figure of merit. Extending it further downstream puts more

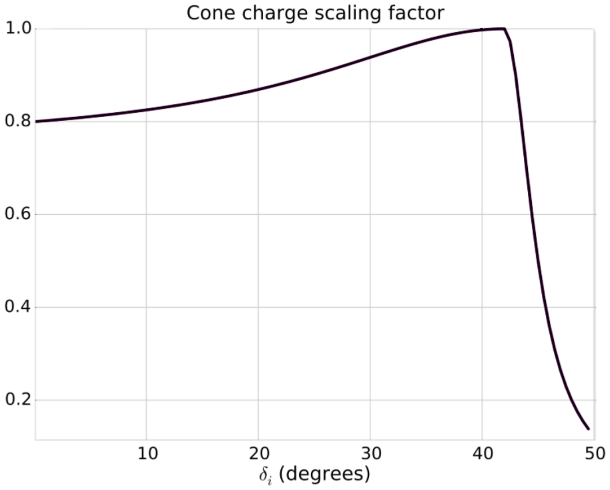

Figure 6.4: Charge weighting used in the cone edge fit figure of merit. Hit charges are weighted such that the figure of merit peaks when the reconstructed cone edge aligns with the true cone edge. Figure by Jingo Wang.

charge outside of the cone, again resulting in a lower weighting and a reduced figure of merit. The maximum figure of merit is achieved when the vertex correctly reproduces the Cherenkov disk edge.

The final figure of merit used for the extended track fit is the sum of Equation 6.2 and Equation 6.6. Its variation for vertices displaced along the transverse and parallel directions is shown in Figure 6.5. The resulting position and angular resolutions are shown in Figure 6.6, giving the results obtained when only PMT hits are used, and when LAPPD hits are also included. The position and angular resolutions achieved by 68% of reconstructed events are marked on Figure 6.6, which evaluate to a position resolution of 12 cm with LAPPDs compared to 38 cm with PMTs alone, and a direction resolution of 5° with LAPPDs compared to 10° with PMTs alone. Reconstruction efficiency is ∼95% for true quasi-elastic events when including both PMTs and LAPPDs. The plot demonstrates the significant improvement in reconstruction capabilities achieved when adding just 5 LAPPDs. This is a result of the improved hit position and time uncertainties, which allow the deduction of a much more tightly peaked time residual distribution. Improved constraint along the track direction tends to require higher, better granularity photocoverage, although specialised techniques may allow events where the disk edge crosses an LAPPD to be much better constrained. This will be part of future reconstruction efforts. The five LAPPDs used for these studies were placed in a '+' pattern on the downstream face of the water tank, as shown in Figure 6.7. These positions were found to offer the best performance based on Monte Carlo studies in which the length of the LAPPDs were artificially extended to span the full height of the tank, and one such elongated LAPPD was placed on each corner of the inner structure. This allowed simulated hits to be recorded over the entire available space that may be covered by the LAPPDs. Vertex resolutions were then compared for various configurations by masking out areas for each LAPPD and only considering hits within these regions.

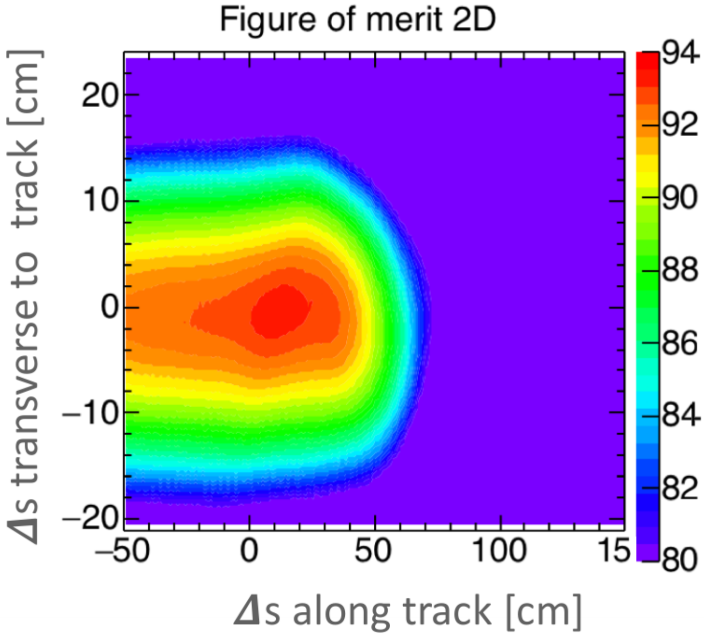

Figure 6.5: Variation in the figure of merit, defined as the sum of Equation 6.2 and Equation 6.6, over a range of vertices displaced along and transverse to the particle direction. Transverse to the track time residuals provide a tight constraint, such that the figure of merit peaks sharply around 0 transverse displacement. Fitting the cone edge provides a much looser constraint, as shown by the broader spread of values, but does also show a peak around 0. Figure by Jingbo Wang.

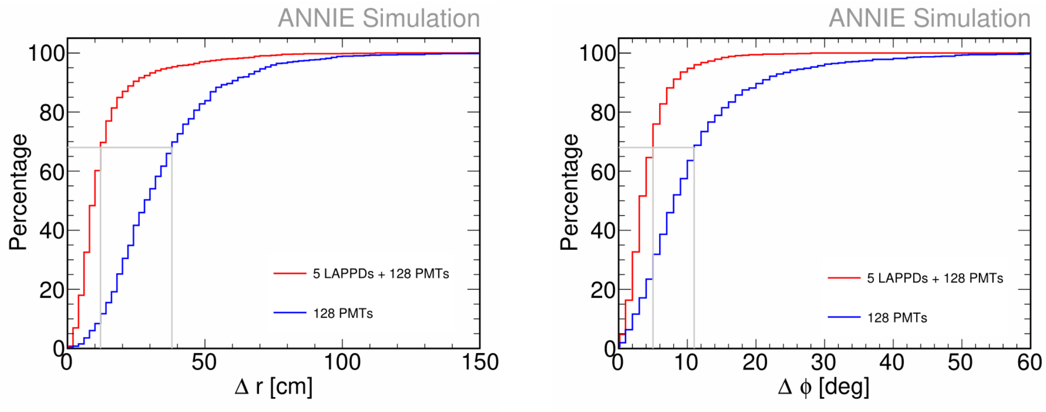

Figure 6.6: Left: Cumulative position resolution when performing reconstruction with PMTs only, or with both PMTs and LAPPDs. With only PMTs 68% of vertices are reconstructed within 38 cm of the true vertex, whereas the inclusion of LAPPDs improves this to 12 cm. Right: Cumulative angular resolution, again with and without the inclusion of LAPPD hits. The improvement in this case is from 10° using PMTs only to 5° with LAPPDs. These results are based on preliminary simulations. Figures by Jingbo Wang.

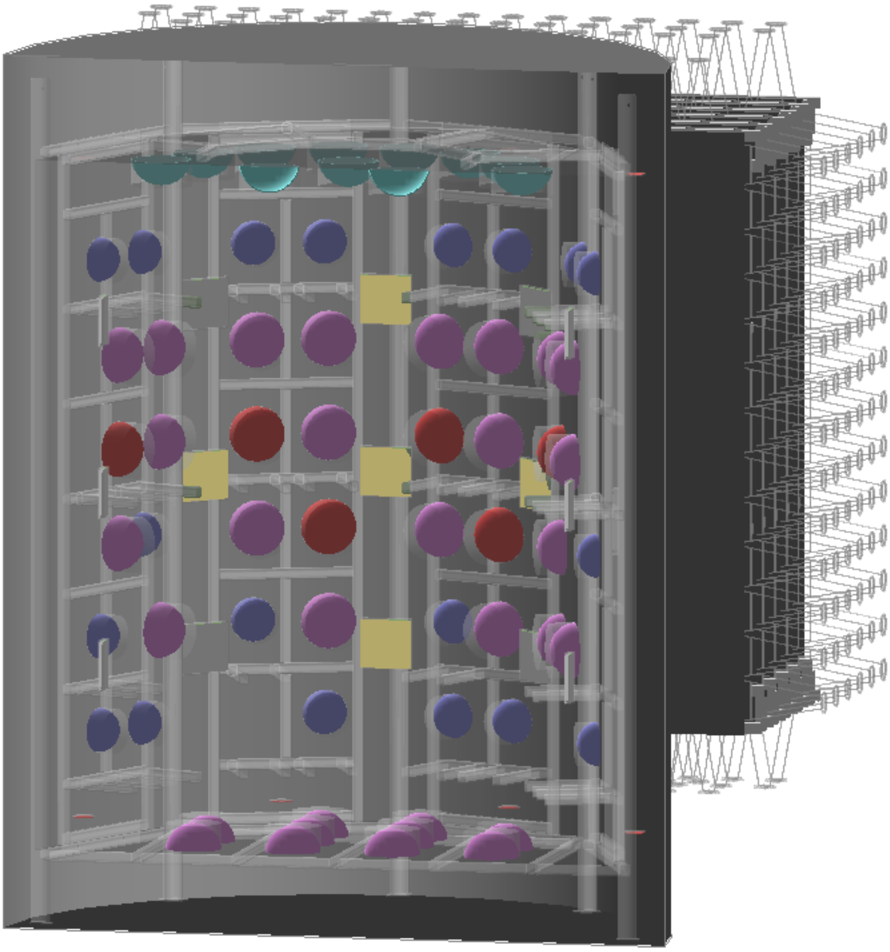

Figure 6.7: Position of LAPPDs used in the vertex resolution studies. The active LAPPDs are highlighted in yellow.

6.4 MRD Track Reconstruction

MRD track reconstruction is based on two separate algorithms, a primary one for long tracks and a secondary one for short stubs with a lower energy threshold.

6.4.1 Long Track Reconstruction

The primary reconstruction algorithm is based on a cellular automation method previously used by the K2K[213] and T2K[214] experiments. Operating within the event time window each MRD paddle may either have a hit or no hit; multiple hits on the same PMT are ignored. Track reconstruction starts by considering the horizontal and vertical views independently. The process is described in the text below, and is shown graphically in Figure 6.8.

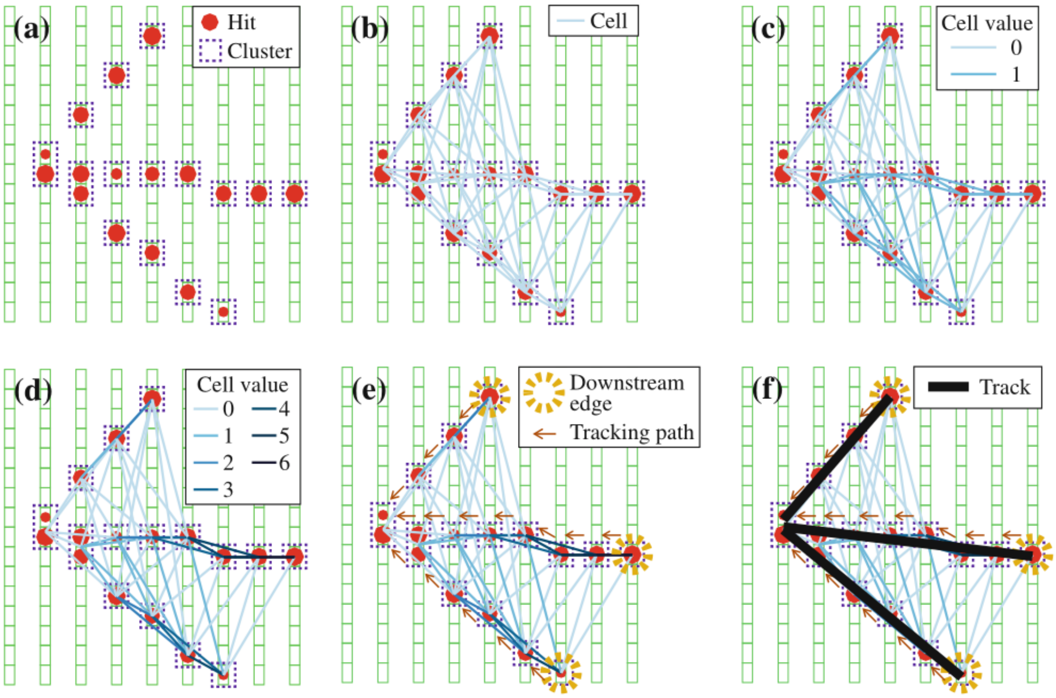

Figure 6.8: Stages of MRD track reconstruction. (a) Hits are combined into clusters, (b) Cells are constructed between clusters, (c)+(d) Cells are paired with neighbours and then each cell status is iterated based on its neighbour statuses to deduce the cell order. (e) Track ends are identified and tracks are constructed. (f) A linear fit determines the track properties. Figure from [214].

1. Within each layer all adjacent hits are grouped into one ‘cluster’.
2. For each cluster, make a ‘cell’ with all clusters in the next two downstream layers. A cell represents a track segment joining two clusters.

3. For each cell, assign an upstream neighbour as any sufficiently aligned cell that shares its upstream cluster. If multiple candidates are avilable, choose the best aligned.
4. Give all cells a starting 'status' of 0. Increment the status of all cells with an upstream neighbour in the same state. Repeat this process until there are no further changes.

Once iteration completes the status of each cell represents the number of upstream neighbours. Cells that are not the upstream neighbour of a cell with a higher state number define the ends of identified tracks. Together with the series of neighbours with sequential state numbers, these cell groups define the 2D tracks.

The next step is to combine the 2D tracks in each view to produce a set of complete 3D tracks. For each possible pairing a figure of merit is calculated, based on:

- The proximity of track start and end points
- Agreement between the two views on which half of the MRD was penetrated[1]
- The steepness of the track from the beam (shallower tracks are preferred)

Tracks in each view are then matched based on the best figure of merit, except in some specific cases. Often at this point there may be fake tracks present, connecting noise hits, forming links between tracks, or joining tracks across a kink. Occasionally these can even have the best figure of merit, despite the presence of a suitable pair of tracks arising from true particles. To disfavour the fakes any candidate pairing involving a track that splits, merges, or joins with other tracks is only made if no alternative is available, even if the alternative has a lower figure of merit. At the end of matching, any remaining unmatched tracks are discarded.

For all surviving tracks a linear fit is performed to all clusters in each view to determine the aggregate properties. The resulting length and angle are used in determining the muon energy loss, and the projected MRD entry point is used to match tracks with those from the tank. An event display showing the track cells and final fit has been previously shown in the calibration section, Figure 3.15.

---

[1] Each layer of the MRD is made up of two halves. Layers with paddles oriented vertically, for example, can still determine which half of the MRD (top or bottom) was hit, allowing a cross-check with the track in horizontally oriented paddles.

6.4.2 Short Track Reconstruction

While the previous technique works well for long tracks, it does have a fairly high energy threshold. Just one cell in each view corresponds to 4 layers of steel being penetrated, or a minimum energy loss of around 500 MeV. To cover lower energies a second algorithm is used. This relies on information from the tank track reconstruction to constrain the region of interest and identify MRD activity that would not be sufficient for standalone reconstruction.

As with long track reconstruction, each view is considered independently. Adjacent hits are combined to form clusters, and all clusters in a front layer are used to start a track. For each track, nearby paddles in the next layer (those within a paddle width on either side of the cluster bounds) are then searched for a cluster, and the closest (if any) is added to the track. From here each track is projected to the next layer along its most recent trajectory, and again a search for nearby clusters is performed. This continues until no suitable clusters are found. A final pruning step then retains only the longest track in any cases where multiple tracks share the same cluster.

No attempt is made to match the tracks from the two MRD planes. These tracks are only considered for matching to a tank track when no suitable longer track is found, and are then matched directly to the tank track independently in each view. The forward projection of the tank track defines a window of acceptance, and any short tracks starting within this window are considered as matching MRD activity.

6.4.3 Combined Results

After combining these two algorithms, preliminary studies suggest that for MRD stopped tracks the track length resolution is $\sim$5 cm at full-width-half-maximum, while the track angle resolution is around 6°, averaged over all reconstructed tracks. The rate of energy loss as a function of angle from the beam axis is determined from a quadratic fit to Monte Carlo data. The muon energy loss is then taken as the product of the track length and the appropriate rate of energy loss. This achieves an energy resolution of around 11%, with a distribution as shown in Figure 6.9.

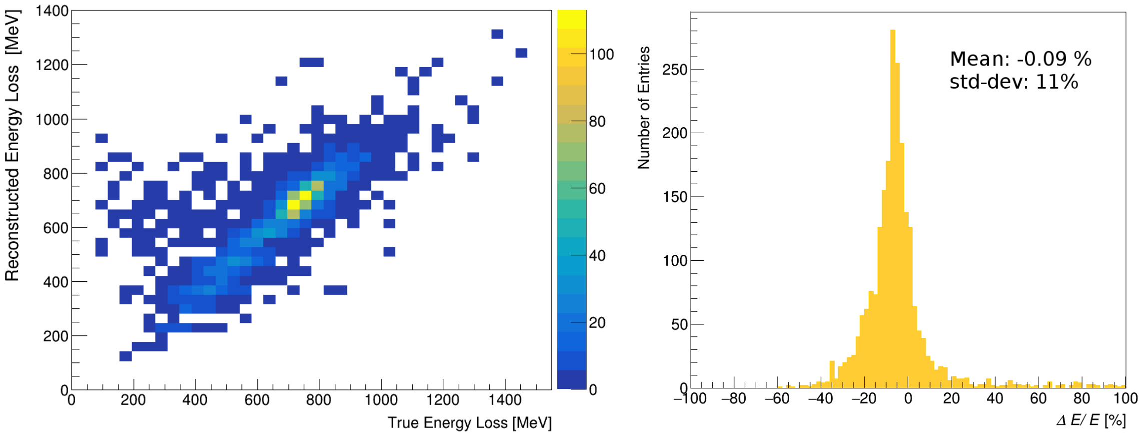

Figure 6.9: Left: Reconstructed vs true energy loss for MRD tracks. Right: Corresponding distribution of relative energy loss errors.

6.5 Tank Energy Reconstruction

As shown by Equation 2.1 the number of photons produced by a charged particle traversing the tank is roughly constant per unit length. For a minimum ionising particle (such as a muon from a neutrino interaction) the total energy loss is proportional to the track length, and therefore also to the number of photons emitted. The detector output, however, must also account for the track geometry, detector geometry, photosensor coverage and charge response. For this ANNIE will be using a machine learning technique developed by collaborators at the University of Edinburgh[153][154].

The reconstruction first uses a Deep Learning Neural Network to generate an accurate estimate of the length of the track in the tank. The network is implemented in the TensorFlow framework[215] and is provided with the following information:

- An initial estimate of the track length, based on the reconstructed vertex and direction
- The total number of hits on all PMTs and LAPPDs
- The times of all hits
- The estimated emission point of the photons that produced each hit, based on the time residual decomposition in Equation 6.1

The performance of the reconstruction is shown in Figure 6.10. The distribution is sharply peaked, with a mean of just 11 cm – considerably better than using the reconstructed vertex alone. These results are based on a representative sample of beam muons, after selecting those that enter the MRD.

The muon energy is then reconstructed using a Boosted Decision Tree with a Gradient Boost[216], implemented using the SciKit Learn library[217]. The tree takes in the following set of variables:

- The reconstructed vertex coordinates
- The distances of the reconstructed vertex from the detector walls
- The length of the track in the tank, calculated with the Deep Neural Network
- The length of the track in the MRD, reconstructed using the methods of section 6.4
- The angle between the reconstructed direction and the beam axis
- The total number of hits on PMTs and LAPPDs

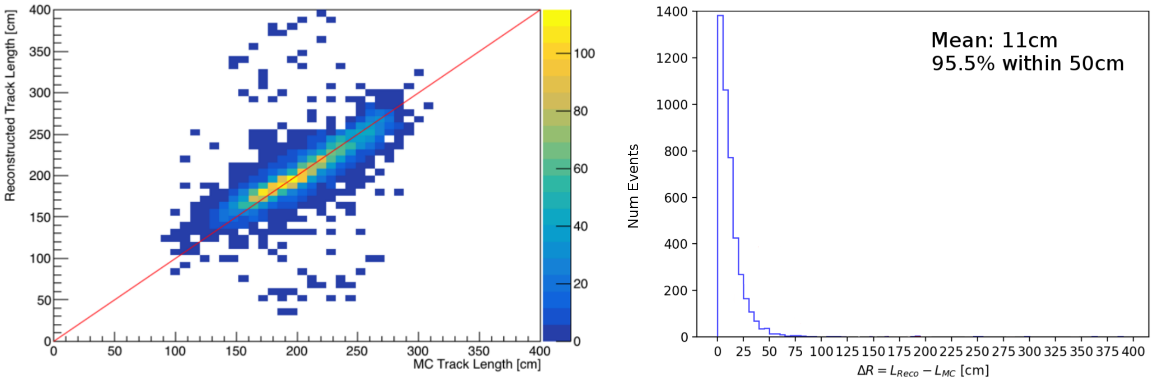

Figure 6.10: Left: True versus reconstructed track length in the tank, for simulated beam muons entering the MRD, calculated using a deep neural network. Right: Corresponding distribution of relative errors.

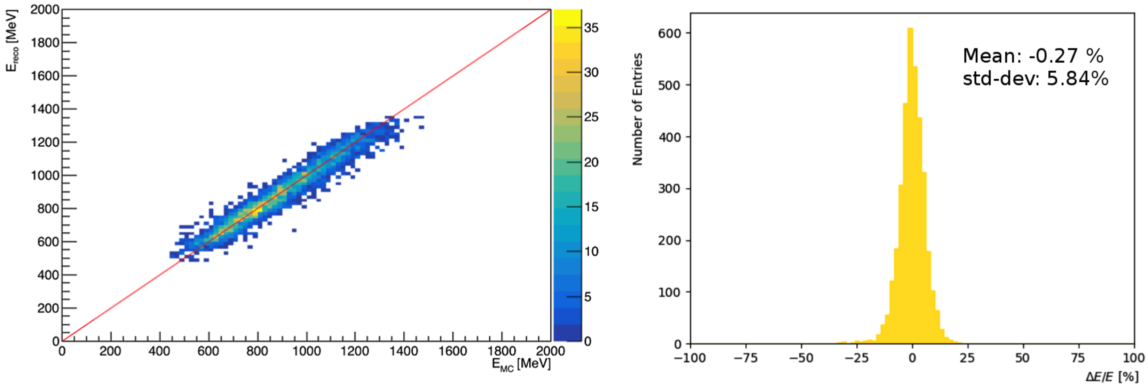

Figure 6.11: Left: True versus reconstructed energy for muons that enter the MRD. The z-axis represents the number of events. Right: Distribution of fractional reconstruction error, defined as $\Delta E = (E_{true} - E_{reco})/E_{true}$.

The model is trained with 1000 boosting stages and a learning rate of 0.01, using a least absolute deviation loss function. These hyperparameters were optimized for the maximum cross-validation using an exhaustive grid search. Results from the reconstruction are shown in Figure 6.11. The muon energy resolution is around 6%.

Given the success of this technique it has also been applied to neutrino energy reconstruction, with initial results showing nearly the same level of success. At the present time the energy loss in the MRD estimated in subsection 6.4.3 is not used, but including this information may help to improve results further. Having reconstructed the muon energy and angle the neutrino energy and momentum transfer are estimated by equations Equation 1.56 and Equation 1.57. Preliminary studies of the accuracy of Q^2 reconstruction are shown in Figure 6.12, but these

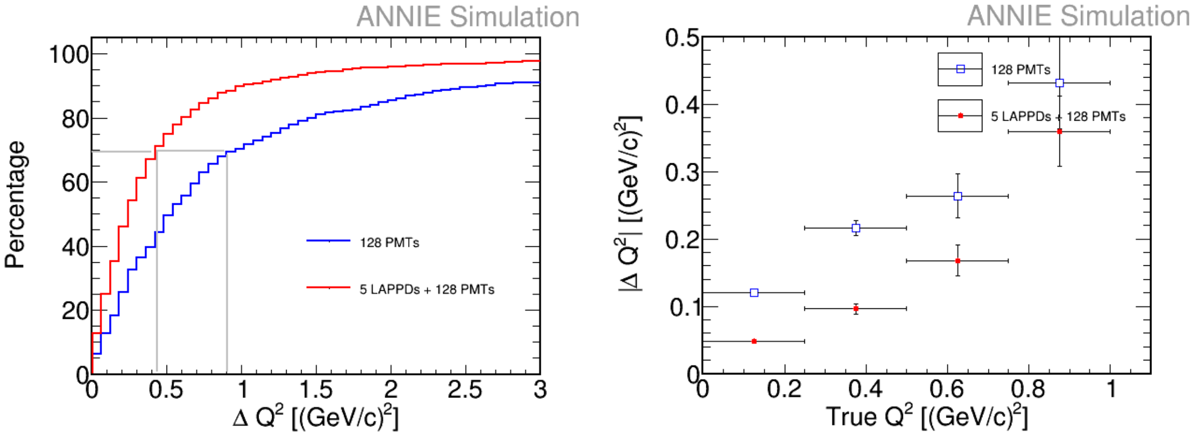

Figure 6.12: Left: Cumulative number of events with a reconstructed squared momentum transfer (Q^2) within a given error. Right: The variation in error as a function of true Q^2. These plots use only true quasi-elastic events selected using Monte Carlo truth information, so are expected to over-estimate reconstruction performance.

results used Monte Carlo truth information to select a pure CCQE sample so are expected to over-estimate performance. Experimental measurements will use the absence of neutrons to select a high purity CCQE sample, but this will not be perfect. Since reconstruction of Q^2 will still rely on this assumption through the use of Equation 1.57 it will be important to keep in mind the implications this has on interpreting measurements, as discussed previously in subsection 1.6.1.

6.6 Event Classification

The first step to classifying event topologies is the identification of multi-ring events, a task conventionally performed using a Hough transform[218]. This involves constructing a vector between the reconstructed vertex and each hit, then projecting a cone about that vector with the Cherenkov angle onto the detector wall. The detector wall is binned into polar angles relative to the reconstructed particle direction, and a count is added to each bin through which the cone perimeter passes. This is repeated for all hits, resulting in a map of the form shown in Figure 6.13. The Hough transform turns rings into peaks, thus turning ring counting into much simpler peak-counting. As each successive ring is identified its associated hits may be subtracted and the algorithm re-run, allowing potentially many overlapping rings to be reconstructed[219]. Unfortunately, given the small size of the ANNIE tank most muons are fully penetrating, producing disks rather than rings and compromising the effectiveness of the method. Some success was seen by introducing an initial edge detection step and performing ring counting

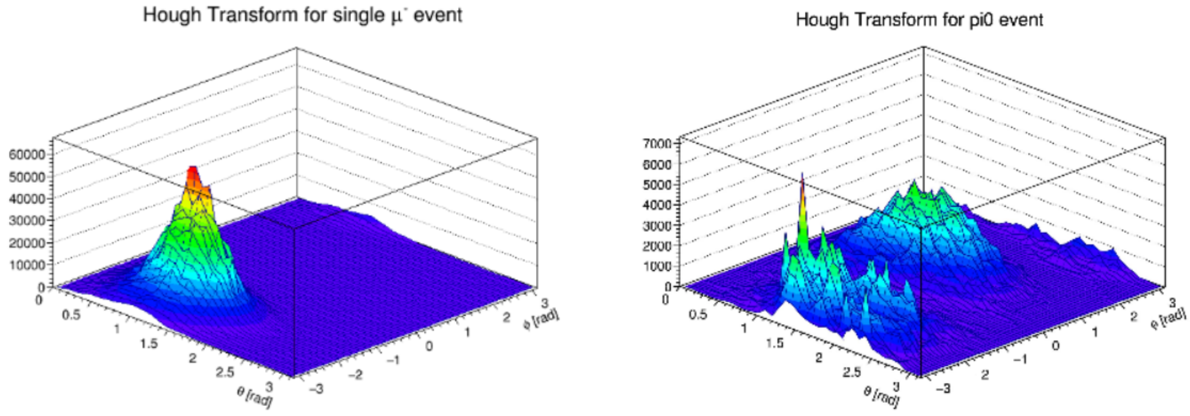

Figure 6.13: Hough transforms of a CCQE interaction (left) and NCπ^0 interaction (right). These images are illustrative of the technique only, and do not represent results from the ANNIE detector. Figure by Matthew Malek.

using only those hits along the identified edges, but further studies are needed to determine the efficacy on real multi-ring events.

The primary limitation of the conventional Hough transform arises from the assumption a point-like vertex. The generalization of this technique to fit extended tracks is the ‘isochron transform’, which divides photon paths into two segments like those used in fitting time residuals (Figure 6.3). These paths are then rotated around the direct line between the track start vertex and the hit location. Doing this with multiple hits produces a maximal overlap at the true emission point of the photon. With many hits and sufficient resolution this method is able to reconstruct complete particle trajectories, including the identification of multiple track events. It would also intrinsically accommodate the Cherenkov discs seen in the ANNIE tank. A demonstration of the potential of this technique, based on simulations from ref [220], is shown in Figure 6.14. The decomposition of photon paths into two components depends on the combination of high timing and spatial resolutions offered by LAPPDs, so this is a powerful method that could be demonstrated for the first time within the ANNIE detector.

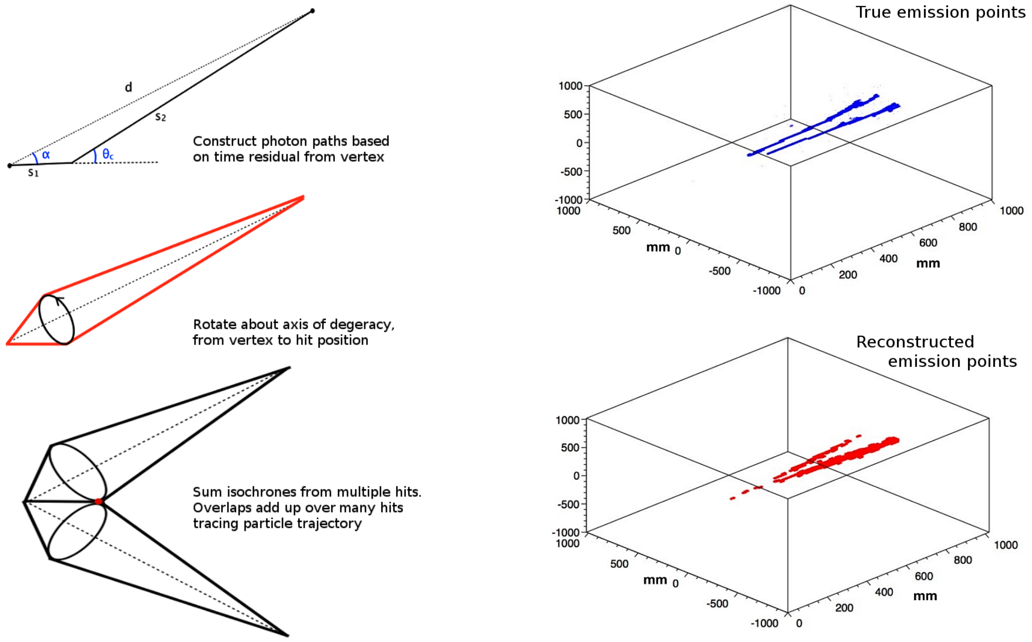

Figure 6.14: Left: Decomposition of the time residual into two parts involves a degeneracy about the line between the vertex and hit position. Rotating about this axis, and combining the result from many hits, allows the complete particle trajectory to be reconstructed. Right: Geant4 simulations of a 1.5 GeV π^0 event, showing the true (top) and reconstructed (bottom) photon emission points in the first two radiation lengths. These are illustrative results based on picosecond detectors in general, but demonstrate techniques that could be developed in the ANNIE detector. Figure from [221].

7

Analysis

7.1 Analysis Strategy

The previous chapter covered the implementational details and performance of the first steps involved in event reconstruction. The analysis of data proceeds fairly linearly through these stages, with cross-section measurements then following from strategy described in section 1.6. Neutron multiplicity measurements will follow a very similar process. So as not to miss the woods for the trees, this section will summarise this strategy, before making some tentative estimates of expected performance.

- A signal from the accelerator opens a prompt window in time with the beam spill. The Central Trigger Card performs the first stage of selection by requiring at least two PMT waveforms exceed a $\sim 4\,\mathrm{mV}$ threshold to trigger an extended readout[1]. Detector data, including tank waveforms, MRD and FMV hits, are written to disk and form the basis of the beam data sample.
- Offline processing begins with event building, which brings together data from each of the sub-detectors to produce a single event for each beam trigger.
- PMT waveforms are baseline subtracted using the pre-beam window, and hits are then identified by scanning each waveform for pulses above a threshold. Hit properties such as charge, peak height and time since last hit are calculated.
- Clusters of hits close in time are identified for both tank and MRD hits, independently.
- Hit cleaning is performed on tank clusters, then vertex reconstruction is performed on the cleaned clusters.

---

[1] the actual voltage thresholds are tuned on a per-PMT basis to account for varying gains

- Track reconstruction is performed on MRD clusters; any tracks found are time-matched to the tank vertices. A tank vertex, together with a coincident MRD track if applicable, defines a sub-event.

- The FMV is searched for hits in coincidence with each sub-event; those with a matching veto hit are flagged as beam-correlated backgrounds.

At this point a readout has been split into potentially many sub-events. Given that a charged-current neutrino interaction is expected in roughly 1 out of every 150 beam spills, the number of readouts with more than one neutrino interaction is expected to be very small. The majority of events are therefore expected contain at most one prompt sub-event without an FMV hit, and potentially multiple delayed sub-events. The analysis from this point onwards will depend on the topology under study, but as an example we can consider CC0π events with an MRD-stopped track and no neutrons in the final state; this is likely to represent the simplest dataset and can be used to perform a CCQE-like cross-section measurement with reduced contamination from non-CCQE events. For such an analysis we identify the primary interaction as the combination of a prompt event within the beam spill window, with no FMV hits, and a coincident MRD track. Analysis then continues as follows:

- The energy of the primary lepton is determined from the prompt sub-event, along with the incident neutrino energy and the momentum transfer under the ν_μ-CCQE assumption.

- The prompt sub-event will also be passed through the event classification algorithms to calculate the likelihood of being the desired event topology. This will be searching for signs of pions in the final state, for example.

- Delayed sub-events will undergo a number of tests to assign likelihoods of being a neutron capture or some other background. Backgrounds include accidental cosmic muons, afterpulsing, nuclear de-excitations and ‘flasher’ events caused by PMT discharge. Likelihoods will be based on observables such as the uniformity of charge distribution, time since the prompt event, energy deposition etc. These tools have yet to be developed, so were not covered in the previous chapter.

At the end of this additional processing the data comprises a set of prompt events, each with its delayed sub-events, fully augmented with reconstructed information and likelihood metrics. A final selection process can then extract a subset of events enriched in the desired topology; for the dataset described above this could include a cut on the likelihood of the prompt event containing pions, together with either no delayed events, or only delayed events with a low likelihood of being neutron captures. The complete set of selection cuts and thresholds will be based on a balance of signal efficiency and sample purity.

The number of events that pass these selections corresponds to N_{int} in Equation 1.63 (this will be an array in kinematic bins for a differential analysis). As discussed in section 1.6, translating this figure into a cross-section requires estimates of the selection process efficiency, the amount of background contamination, the smearing matrix, and the incident neutrino flux. The flux is obtained as described in section 5.2; all remaining quantities must be carefully estimated based on Monte Carlo studies with, wherever possible, data-driven validation.

7.2 Uncertainties

7.2.1 Statistical Errors

Its close proximity to the BNB provides ANNIE with a high rate of interactions that will minimize statistical errors. A list of preliminary selection efficiencies for each of the steps in the analysis process discussed above is given in Table 7.1. Given an expected rate of ∼363,000 ν_μ charged-current events per year in the entire ANNIE tank and a combined selection and reconstruction efficiency of ∼3% we can expect to obtain ∼11,500 events per year in the dataset for this particular topology, corresponding to a standard error on the integrated cross-section on the order of ∼1%.

Source	**Efficiency** [%]	Note
Readout DeadTime	95	From VME implementation
FMV Cut	100	No in-time FMV hits
MRD Track Cut	48	Efficiency of reconstruction success
Tank Vertex Cut	80	Efficiency of fit success
Fiducial Volume Cut	9	Based on 2.5 ton fiducial volume
Event Classifier Cut	95	Based on 5 LAPPDs
No Neutrons Cut	95	Tentative Estimate
Total	3	Integrated over final state kinematics

Table 7.1: Efficiencies of each step in a proposed reduction chain for a sample of CC0π events with an MRD-stopped track and no neutrons in the final state. Efficiencies are integrated over final state kinematics where applicable.

7.2.2 Systematic Errors

For the majority of accelerator neutrino experiments the greater contribution to total uncertainty arises from systematic errors. The dominant source of systematic error, which has already been discussed in detail in section 5.2, is that of the neutrino flux. Still, the data reduction process introduces many more sources of error into the measurement and these must also be accounted for. Each cut in the data reduction process has some associated efficiency for capturing desired (signal) events; if this efficiency is known perfectly one can convert, within statistical precision, the number of signal events after selection to the number before selection. On the other hand any uncertainty on the efficiency results in uncertainty on the true number of signal events before selection. Similarly each selection will accept some undesired (background) events. If the rate of background events before selection is known, and the fraction that pass selection is known, then the number of accepted background events is also known. This can then be subtracted from the total number of accepted events to obtain the number of accepted signal events. Any uncertainty on either the rate of background events or the background acceptance will contribute to uncertainty on the number of signal events and the derived cross-section.
Such uncertainties, on signal efficiency, true background rate and background acceptance, must be determined for all reduction steps. A summary of the main contributions is given in Table 7.2.

Source	Contributes To
Missed FMV Hits	Background
FMV Accidental Mis-tag	Efficiency
Formation of Prompt Tank Cluster from True Event	Efficiency
Formation of Prompt Tank Cluster from Accidental Event	Background
Successful Reconstruction of MRD Track	Efficiency
Accidentals or Noise Reconstructed as MRD Track	Background
Successful Reconstruction of Prompt Tank Vertex	Efficiency
Correctly Tag Event in Fiducial Volume	Efficiency
Mis-tag Event as in Fiducial Volume	Background
Correctly Tag Event as CC0π	Efficiency
Mis-tag Event as CC0π	Background
Mis-tag Delayed Event as Neutron Capture	Efficiency
Missed Tank Cluster from True Neutron Capture	Background
Mis-tag Delayed Event as Not Neutron Capture	Background

Table 7.2: Sources of systematic error on a CC0π cross-section measurement.

7.2.3 Systematic Sources

The efficiency of the FMV for accepting signal events (with no penetrating muon) depends on the likelihood of flagging an accidental FMV hit in coincidence with a tank event, thus falsely rejecting the tank event. Owing to the low rate of neutrino interactions and the tight timing cut on the coincidence requirement, the rate of such cases is expected to be negligible; the efficiency of the FMV for accepting signal events can be considered 100% with negligible error.

Considering background acceptance, the FMV has a $72 \pm 8\,\%$ efficiency for flagging dirt interactions via the entering muon (Figure 3.17). This indicates that there is an an $8\,\%$ uncertainty on the fraction of background events accepted. The number of events this corresponds to depends on the rate of upstream dirt interactions. Of course, we cannot estimate this rate based on the flux and cross-section (which would be circular in a measurement of cross-section) but we may instead use a data-driven method. The background contribution of in-tank events to dirt events is negligible (due to the low rate of in-tank events and tight FMV timing constraint), so we can use the measured rate of dirt interactions together with the FMV detection efficiency to determine the true rate of dirt interactions: $N_{dirt} = N_{tagged}/\epsilon_{FMV}$. We can then invert the efficiency to obtain the rate of dirt interactions without an FMV tag: $N_{untagged} = N_{dirt} * (1 - \epsilon_{FMV}) = N_{tagged} * \frac{1-\epsilon_{FMV}}{\epsilon_{FMV}}$. The final uncertainty on the rate of untagged dirt backgrounds then comes from uncertainty on the FMV efficiency (8%) and the statistical uncertainty on the measured rate of dirt events ($\mathcal{O}(\sqrt{N_{dirt}}) \sim \sqrt{78,000\,yr^{-1}} \sim 0.3\%$.

The requirement of at least 10 hits within 50 ns for the formation of a tank cluster places a low energy threshold on any tank activity. For the prompt event the requirement of a coincident MRD track will independently eliminate essentially all background sources other than untagged dirt muons, while also placing a much higher threshold on the muon energy and thus the amount of tank activity. For both signal and backgrounds, therefore, the efficiency for tank cluster creation is likely to be close to 100%, with negligible error.

The efficiency of MRD track reconstruction strongly depends on the length of the track, as shown in Figure 7.1 left. This plot is derived from a simulated set of beam muons, with track length taken as the difference between the point of MRD entry and stopping or exit. The distribution of MRD track lengths for all events, successfully reconstructed or otherwise, is given in Figure 7.1 right. The main source of uncertainty is this estimated track length distribution; factors such as the paddle efficiencies or rate and distribution of noise hits can be minimized with proper calibration of the Monte Carlo simulations. In order to quantify the uncertainty arising from the track length distribution, the simulated beam flux will need to be varied according to its uncertainty and the resulting variation in MRD track lengths determined. The contribution to the uncertainty in selection efficiency is then obtained from the variation in

MRD track reconstruction efficiencies for each of these distributions. As muon kinematics from charged-current interactions are relatively well understood this systematic can be expected to be small.

Background contributions from ‘fake’ MRD tracks could occur if dark noise hits, cosmic muons, or otherwise unrelated MRD activity is mistakenly matched to a prompt tank event. The likelihood of creating ‘high energy’ fake tracks is small, but as discussed in subsection 6.4.2 only minimal MRD activity is required to count as an MRD track in order to push the energy threshold as low as possible. The trade-off is an increased chance of erroneous reconstructions. Nonetheless, the expected contribution is expected to be small, and could be studied by applying the MRD track reconstruction to random windows in the off-beam region, combined with ‘injected’ tank cluster information from unrelated beam events.

Reconstruction of tank vertices is not always successful as sometimes the minimization procedure does not converge. For a sample of beam CCQE events with a muon entering the MRD, around 80% of fits converged when using hits just from PMTs, rising to 95% when adding in 5 LAPPDs. Convergence may fail if there are too few hits, excessive noise or if the fitter is given an unsuitable seed. Studies on the relationship between convergence and event kinematics will be needed to establish uncertainties on the vertex reconstruction efficiency, and the resulting impact on the measured cross-sections.

The efficiency of reconstruction, both of prompt and delayed events, will vary with true vertex as discussed in subsection 2.11.1; a fiducial volume cut will be applied to limit this variation. The acceptance of the fiducial volume cut will depend on the enclosed volume and on the efficiency of fit convergence at all points within it. These can be accounted for by integrating the convergence efficiency over the fiducial volume, defined with a perimeter varied over a span defined by the vertex resolution. It is likely that this cut will also have some dependence on the event topology; reconstruction of mis-tagged CC1π events, for example, may have a different acceptance to true CC0π events due to differences in the distribution of light. This will result in a correlation between the background acceptance of the fiducial volume cut and that of the event classifier.

Many of the event classification algorithms being developed by ANNIE are based on machine learning techniques. These output a set of likelihoods that the event belongs to each topological class. Current tools using PMT hits achieve an 85% efficiency for correctly identifying CC0π tank events, with 17% of both CCNπ and NCQE events being falsely labelled as CC0π. With the inclusion of LAPPD hits these figures could see substantial improvement, but taking full advantage of the additional data will require new techniques to be developed. Being acutely dependent on the details of final state products, uncertainties in these efficiencies are driven

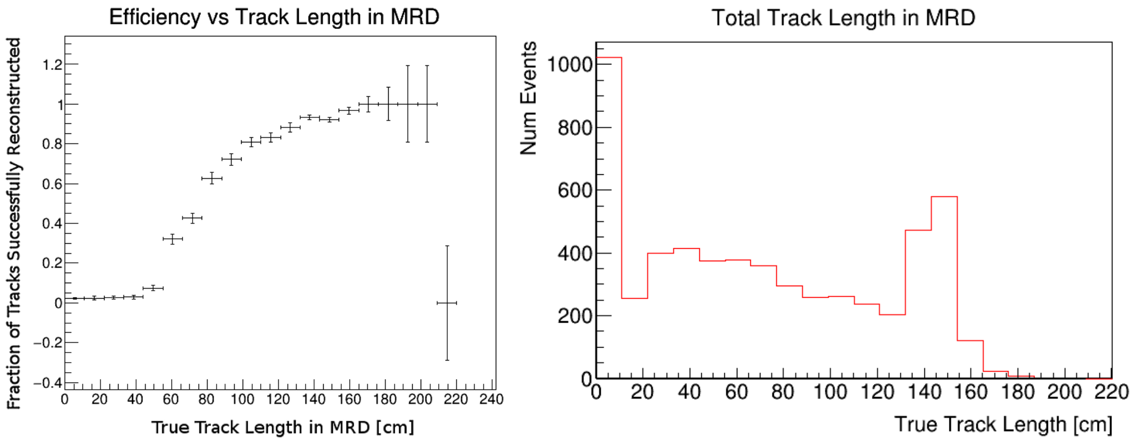

Figure 7.1: Left: Efficiency of MRD track reconstruction, as a function of true length of the muon in the MRD. Right: Distribution of track lengths in the MRD, derived from a simulation of beam muons.

by those on the neutrino interaction models and the relative cross-sections for each interaction channel. Since these are poorly understood extensive studies will be particularly important to establish an appropriate error budget. Such studies generally involve varying the input parameters of the neutrino simulations and assessing the impact on classifier performance. These algorithms are still under development and performance can vary quite widely, so the associated uncertainties are difficult to estimate at this time.

The same principles regarding vertex reconstruction efficiency, event classification and fiducial volume acceptance also apply to identifying neutrons in delayed events. In this case the underlying uncertainties arise from accurate modelling of afterpulse characteristics, the rate of nuclear de-excitations and the rate of flasher events. The methods to be used for this classification have yet to be determined, but captures on gadolinium should yield sufficient visible energy to make their identification high efficiency with low uncertainty.

Contributions to the smearing matrix will come both from the resolutions on the corresponding kinematic variables (muon energy and angle errors), but also from distortion of the measured spectrum introduced by background contributions. This is because backgrounds are not likely to have the same cross-sectional energy-dependence as the desired topology, so their addition to the data sample will change the distribution as well as the normalization of the extracted cross-section. These errors are again estimated with Monte Carlo studies, but may also use background-enhanced side-band samples to provide data-driven constraints.

8

The Future

8.1 Next Steps

As ANNIE undergoes commissioning over the coming months the focus will be on tuning the models to achieve the best agreement with the observed detector response. The previous chapter has shown some of the early work being conducted, which so far show good agreement between simulation and data. Further work will fine tune the light yield, neutron capture efficiencies and timing characteristics. Critically this calibration will determine the uncertainties associated with future physics measurements. Physics quality data taking is expected to begin in 2021.

At the time of writing the tank commissioning does not include LAPPDs. All of ANNIE's first batch of 5 LAPPDs are in-hand and undergoing characterisation, but further work is needed on ex situ calibration before they will be deployed. Reconstruction efforts will crucially rely on the ability to identify multi-photon hits. Previously published results[168] demonstrated a photon identification efficiency of up to 74% for as many as 25 hits on one LAPPD, rising to 95% efficiency for 5 hits. This work will need to be integrated into the ToolAnalysis framework for future holistic reconstruction. Once this foundation is in place more advanced algorithms for track reconstruction and event classification can be worked on, fully leveraging the capabilities of the LAPPDs.

Cross-section measurements may be some of the first physics results produced by the experiment. Nearly all the elements required for an inclusive charged current cross-section measurement are in place, with many of the associated uncertainties being determined as part of the calibration effort. Exclusive studies will rely on the LAPPDs for event classification and complex topological reconstruction. Investigations into potential methods are already underway.

In the meantime, further analysis of the skyshine neutron background, measured during the previous data-taking run, may be of use to other experiments on the Booster Neutrino Beamline. Since the neutron stopping power of water is well known, and the skyshine neutron flux has been measured at multiple depths within the detector, the energy spectrum of skyshine neutrons can be determined from the rate of flux reduction. An analysis of this spectrum is currently ongoing.

8.2 Additional Opportunities

8.2.1 Water-based Liquid Scintillator

Looking further ahead there are more opportunities to leverage the ANNIE detector as a testbed for new technologies. Plans are underway to trial 400 litres of Water-based Liquid Scintillator (WbLS)[222], using a similar vessel as in the neutron background measurement but without the optical isolation. Liquid scintillators provide around 100 times the light yield of Cherenkov emission alone, substantially lowering the threshold for detection and enabling a calorimetric component of energy reconstruction. While experiments such as MiniBooNE have made use of both light components, the transparency of organic scintillator is much lower than that of pure water, and the Cherenkov component would be absorbed over the large scales possible with pure water detectors. Water-based liquid scintillators use surfactants to enable the suspension of small drops of scintillator oil, known as micelles, within pure water (Figure 8.1). The fractional light yield is close to the concentration, but with transparencies much closer to that of pure water[223]. Moreover the ionic nature of the solvent allows doping with metals such as gadolinium.

The emission of scintillation light is delayed by a relaxation time of ∼5–20 ns, compared to prompt Cherenkov emission that typically spans ∼0–5 ns for a point-like event. The timing resolution offered by LAPPDs allows the separation of these two components of light, allowing maximal information to be extracted about the event[224][225][226]. This technique is a central component of the proposed next-generation Theia detector[227], and ANNIE could provide its first demonstration in a neutrino experiment.

8.2.2 Beam Spill Time Slicing

The wave of neutrinos arriving from an accelerator intrinsically carries a time-dependent energy spread. This phenomenon arises from correlations between the parent pion energy, the energy of the neutrino produced in its decay, and the distance the parent travels before decaying. Lower

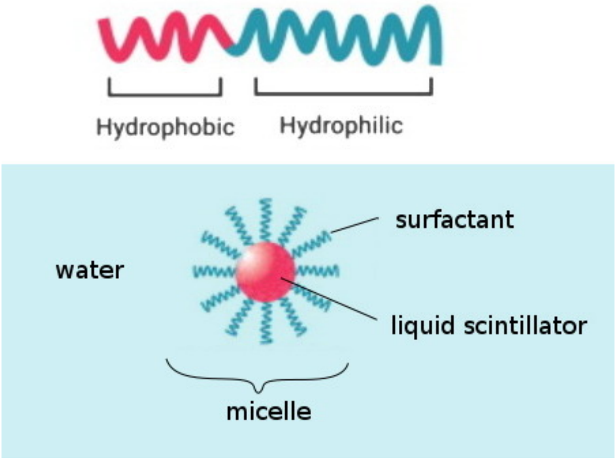

Figure 8.1: Schematic of a micelle of liquid scintillator oil suspended in water. Figure adapted from [228].

energy hadrons have a lower velocity and decay further upstream than more energetic ones, even after accounting for time dilation of the decay lifetime. Since all neutrinos propagate at essentially the speed of light this additional distance means lower energy neutrinos have further to travel, arriving at the detector up to ∼1ns later. With sufficient timing resolution and short enough beam spills different energy components can be stroboscopically separated[229]. Observation of this effect would require a rebunching of the BNB to narrow the width of the incident proton bunches. Commercially available superconducting RF cavities could reduce the bunch widths to ∼200 ps, enabling such separation while preserving the total number of protons per bunch. The physics potential and feasibility of implementation was recently discussed at a Fermilab workshop[230], and if implemented could be observed by the ANNIE detector.

9

Conclusion

This thesis has discussed the ANNIE detector, its goals, its construction and the current state of event reconstruction. ANNIE is poised to measure charged-current neutrino cross-sections and neutron multiplicities for interactions on water in the energy range of 0.5–3 GeV. Neutrino cross-sections on oxygen are at the present time poorly measured[107][108][109][110]. The results from ANNIE will provide crucial data for cross-section modelling, and direct measurements that can be used in present and future water target experiments without requiring model-dependent extrapolation. The presence of neutrons in charged-current neutrino events is a strong indication of inelasticity. This can be used to isolate a purer sample of quasi-elastic events for improved energy reconstruction, and will be a critical handle for signal-background separation in searches for proton decay and supernova relic neutrinos. Understanding the expected neutron yield from neutrino interactions is an essential pre-requisite to applying these methods. Current models show large discrepancies when compared with observed proton multiplicity, and while no data yet exists the same inaccuracies are expected for neutrons. The measurements provided by ANNIE will be crucial input for improving these models, and as direct input for the next generation of detectors utilising neutron detection.

Beyond its physics goals ANNIE will demonstrate the first use of Large Area Picosecond Photodetectors (LAPPDs) in a high energy physics experiment. These detectors provide simultaneous picosecond timing and centimetre spatial resolution, enabling numerous new techniques in event reconstruction. ANNIE offers a testbed to develop and demonstrate these new methods, and they will be essential to realising its physics goals.

At the current time ANNIE is undergoing commissioning, with hopes to begin taking physics quality data at the start of the next beam season. First physics data is expected in 2021.

References

[1] E. Noether. **Invariant variation problems**. *Transport Theory and Statistical Physics*, **1**(3):186–207, Jan 1971. doi:10.1080/00411457108231446. 1

[2] A. collaboration. **Summary of new ATLAS results for ICHEP 2020**. Accessed: 2020-10-14, July 2020. 1

[3] R. Carlin. **CMS Highlights**. Talk at the 2020 International Conference on High-Energy Physics (ICHEP), August 2020. 1

[4] D. Akerib et al. **Results from a search for dark matter in the complete LUX exposure**. *Phys. Rev. Lett.*, **118**(2):021303, 2017. doi:10.1103/PhysRevLett.118.021303. 1

[5] E. Aprile et al. **Light Dark Matter Search with Ionization Signals in XENON1T**. *Phys. Rev. Lett.*, **123**:251801, Dec 2019. doi:10.1103/PhysRevLett.123.251801. 1

[6] M. Miura. **Search for Nucleon Decay in Super-Kamiokande**. *Nuclear and Particle Physics Proceedings*, **273-275**:516 – 521, 2016. 37th International Conference on High Energy Physics (ICHEP). doi:https://doi.org/10.1016/j.nuclphysbps.2015.09.076. 1

[7] B. Odom, D. Hanneke, B. D'Urso, and G. Gabrielse. **New Measurement of the Electron Magnetic Moment Using a One-Electron Quantum Cyclotron**. *Physical review letters*, **97**:030801, 08 2006. doi:10.1103/PHYSREVLETT.97.030801. 1

[8] T. Humair. **General Lepton Universality Tests**. Talk at the 2019 Standard Model at the LHC Conference, April 2019. 1

[9] T. Gorringe and D. Hertzog. **Precision muon physics**. *Progress in Particle and Nuclear Physics*, **84**:73 – 123, 2015. doi:https://doi.org/10.1016/j.ppnp.2015.06.001. 1

[10] G. Pásztor. **Precision tests of the Standard Model at the LHC with the ATLAS and CMS detectors**. *PoS*, **FFK2019**:005, 2020. doi:10.22323/1.353.0005. 1

[11] A. Joyce, B. Jain, J. Khoury, and M. Trodden. **Beyond the Cosmological Standard Model**. *Phys. Rept.*, **568**:1–98, 2015. doi:10.1016/j.physrep.2014.12.002. 2

[12] J. Solà. **Cosmological constant and vacuum energy: old and new ideas**. *Journal of Physics: Conference Series*, **453**:012015, aug 2013. doi:10.1088/1742-6596/453/1/012015. 2

[13] S. D. M. White, C. S. Frenk, and M. Davis. **Clustering in a neutrino-dominated universe**. *The Astrophysical Journal*, **274**:L1–L5, November 1983. doi:10.1086/184139. 2

[14] M. Tanabashi et al. (Particle Data Group). **Review of Particle Physics**. *Physics Review D*, **98**:030001, 2018. 2019 update for the 2018 edition. doi:10.1103/PhysRevD.98.030001. 2, 3

[15] W. Bernreuther. **CP violation and baryogenesis**. *Lect. Notes Phys.*, **591**(hep-ph/0205279. PITHA-2002-08):237–293. 57 p, May 2002. 2

[16] T. Kajita, E. Kearns, and M. Shiozawa. **Establishing atmospheric neutrino oscillations with Super-Kamiokande**. *Nuclear Physics B*, **908**:14 – 29, 2016. Neutrino Oscillations: Celebrating the Nobel Prize in Physics 2015. doi:https://doi.org/10.1016/j.nuclphysb.2016.04.017. 3

[17] M. Aker et al. **Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN**. *Phys. Rev. Lett.*, **123**(22):221802, 2019. doi:10.1103/PhysRevLett.123.221802. 3

[18] S. Bharti, S. Prakash, U. Rahaman, and S. Uma Sankar. **Understanding the degeneracies in NOνA data**. *JHEP*, **09**:036, 2018. doi:10.1007/JHEP09(2018)036. 3

[19] W. Buchmuller, P. Di Bari, and M. Plumacher. **Leptogenesis for pedestrians**. *Annals Phys.*, **315**:305–351, 2005. doi:10.1016/j.aop.2004.02.003. 5

[20] C. Giunti, C. W. Kim, and U. W. Lee. **When do neutrinos really oscillate? Quantum mechanics of neutrino oscillations**. *Phys. Rev. D*, **44**:3635–3640, Dec 1991. doi:10.1103/PhysRevD.44.3635. 13

[21] S. King. **Discrete Symmetries and Models of Flavour Mixing**. *Journal of Physics: Conference Series*, **631**:012005, 07 2015. doi:10.1088/1742-6596/631/1/012005. 14

[22] D. H. Wilkinson. **Limits to second-class nucleonic and mesonic currents**. *Eur. Phys. J.*, **A7**:307–315, 2000. 15

[23] J. Arrington, K. de Jager, and C. F. Perdrisat. **Nucleon Form Factors – A Jefferson Lab Perspective**. *Journal of Physics: Conference Series*, **299**:012002, may 2011. doi:10.1088/1742-6596/299/1/012002. 15

[24] S. L. Adler. **Tests of the Conserved Vector Current and Partially Conserved Axial-Vector Current Hypotheses in High-Energy Neutrino Reactions**. *Phys. Rev.*, **135**:B963–B966, Aug 1964. doi:10.1103/PhysRev.135.B963. 15

[25] M. Gell-Mann and M. Levy. **The axial vector current in beta decay**. *Nuovo Cim.*, **16**:705, 1960. doi:10.1007/BF02859738. 16

[26] A. S. Meyer, M. Betancourt, R. Gran, and R. J. Hill. **Deuterium target data for precision neutrino-nucleus cross sections**. *Phys. Rev.*, **D93**(11):113015, 2016. doi:10.1103/PhysRevD.93.113015. 16

[27] B. Bhattacharya, G. Paz, and A. J. Tropiano. **Model-independent determination of the axial mass parameter in quasielastic antineutrino-nucleon scattering**. *Phys. Rev. D*, **92**(11):113011, 2015. doi:10.1103/PhysRevD.92.113011. 16

[28] C. Alexandrou, M. Constantinou, K. Hadjiyiannakou, K. Jansen, C. Kallidonis, G. Koutsou, and A. Vaquero Aviles-Casco. **Nucleon axial form factors using $N_f = 2$ twisted mass fermions with a physical value of the pion mass**. *Phys. Rev.*, **D96**(5):054507, 2017. doi:10.1103/PhysRevD.96.054507. 16

[29] T. Leitner, O. Buss, L. Alvarez-Ruso, and U. Mosel. **Electron- and neutrino-nucleus scattering from the quasielastic to the resonance region**. *Phys. Rev.*, **C79**:034601, 2009. doi:10.1103/PhysRevC.79.034601. 17

[30] J. Zmuda and K. Graczyk. **Electroweak single pion production and form factors of the Δ(1232) resonance**. *AIP Conf. Proc.*, **1680**(1):020013, 2015. doi:10.1063/1.4931872. 18

[31] J. M. Conrad, M. H. Shaevitz, and T. Bolton. **Precision measurements with high-energy neutrino beams**. *Rev. Mod. Phys.*, **70**:1341–1392, 1998. doi:10.1103/RevModPhys.70.1341. 18

[32] A. V. Manohar. **An Introduction to spin dependent deep inelastic scattering**. In *Lake Louise Winter Institute: Symmetry and Spin in the Standard Model Lake Louise, Alberta, Canada, February 23-29, 1992*, pages 1–46, 1992. 18

[33] A. Bodek and U. Yang. **Unified approach for modelling neutrino and electron nucleon scattering cross sections from very high Q**2 to Q**2 = 0**. *AIP Conf. Proc.*, **721**(1):358–362, 2004. doi:10.1063/1.1818434. 19

[34] A. Bodek and U. Yang. **NUFACT09 update to the Bodek-Yang Unified Model for Electron- and Neutrino- Nucleon Scattering Cross Sections**. *AIP Conference Proceedings*, **1222**(1):233–237, 2010. doi:10.1063/1.3399303. 19

[35] C. H. Llewellyn Smith. **Neutrino Reactions at Accelerator Energies**. *Phys. Rept.*, **3**:261–379, 1972. doi:10.1016/0370-1573(72)90010-5. 21

[36] J. G. Morfin, J. Nieves, and J. T. Sobczyk. **Recent Developments in Neutrino/Antineutrino - Nucleus Interactions**. *Adv. High Energy Phys.*, **2012**:934597, 2012. doi:10.1155/2012/934597. 23, 35

[37] L. Alvarez-Ruso, Y. Hayato, and J. Nieves. **Progress and open questions in the physics of neutrino cross sections at intermediate energies**. *New J. Phys.*, **16**:075015, 2014. doi:10.1088/1367-2630/16/7/075015. 23, 35

[38] D. Rein and L. M. Sehgal. **Neutrino Excitation of Baryon Resonances and Single Pion Production**. *Annals Phys.*, **133**:79–153, 1981. doi:10.1016/0003-4916(81)90242-6. 23

[39] T. Leitner, O. Buss, U. Mosel, and L. Alvarez-Ruso. **Neutrino scattering with nuclei: Theory of low energy nuclear effects and its applications**. *PoS*, **NUFACT08**:009, 2008. doi:10.22323/1.074.0009. 23, 34

[40] K. M. Graczyk and J. T. Sobczyk. **Form Factors in the Quark Resonance Model**. *Phys. Rev.*, **D77**:053001, 2008. [Erratum: Phys. Rev.D79,079903(2009)]. doi:10.1103/PhysRevD.79.079903,10.1103/PhysRevD.77.053001. 24

[41] S. X. Nakamura, H. Kamano, T. S. H. Lee, and T. Sato. **Neutrino-induced meson productions off nucleon at forward limit in nucleon resonance region**. *AIP Conf. Proc.*, **1663**(1):070005, 2015. doi:10.1063/1.4919486. 24

[42] A. A. Aguilar-Arevalo et al. **Measurement of Neutrino-Induced Charged-Current Charged Pion Production Cross Sections on Mineral Oil at $E_\nu \sim 1$ GeV**. *Phys. Rev.*, **D83**:052007, 2011. doi:10.1103/PhysRevD.83.052007. 24

[43] S. C. Pieper, R. B. Wiringa, and V. R. Pandharipande. **Variational calculation of the ground state of ^{16}O**. *Phys. Rev. C*, **46**:1741–1756, Nov 1992. doi:10.1103/PhysRevC.46.1741. 26

[44] J. Carlson, S. Gandolfi, F. Pederiva, S. C. Pieper, R. Schiavilla, K. E. Schmidt, and R. B. Wiringa. **Quantum Monte Carlo methods for nuclear physics**. *Rev. Mod. Phys.*, **87**:1067, 2015. doi:10.1103/RevModPhys.87.1067. 26

[45] R. Smith and E. Moniz. **Neutrino reactions on nuclear targets**. *Nuclear Physics B*, **43**:605 – 622, 1972. [Erratum: Nucl. Phys. B43,605(1972)]. doi:https://doi.org/10.1016/0550-3213(72)90040-5. 26

[46] A. M. Ankowski, O. Benhar, and M. Sakuda. **Improving the accuracy of neutrino energy reconstruction in charged-current quasielastic scattering off nuclear targets**. *Phys. Rev.*, **D91**(3):033005, 2015. doi:10.1103/PhysRevD.91.033005. 26

[47] S. Singh and E. Oset. **Quasielastic neutrino (antineutrino) reactions in nuclei and the axial-vector form factor of the nucleon**. *Nuclear Physics A*, **542**(4):587 – 615, 1992. doi:https://doi.org/10.1016/0375-9474(92)90259-M. 27

[48] M. Gell-Mann and K. A. Brueckner. **Correlation Energy of an Electron Gas at High Density**. *Phys. Rev.*, **106**:364–368, Apr 1957. doi:10.1103/PhysRev.106.364. 27

[49] J. Nieves, J. E. Amaro, and M. Valverde. **Inclusive quasielastic neutrino reactions**. *Phys. Rev.*, **C70**:055503, 2004. [Erratum: Phys. Rev.C72,019902(2005)]. doi:10.1103/PhysRevC.70.055503,10.1103/PhysRevC.72.019902. 27, 34

[50] J. Nieves, I. Ruiz Simo, and M. J. Vicente Vacas. **Inclusive Charged–Current Neutrino–Nucleus Reactions**. *Phys. Rev.*, **C83**:045501, 2011. doi:10.1103/PhysRevC.83.045501. 27, 31, 33

[51] L. Alvarez-Ruso, O. Buss, T. Leitner, and U. Mosel. **Quasielastic Scattering at MiniBooNE Energies**. *AIP Conf. Proc.*, **1189**(1):151–156, 2009. doi:10.1063/1.3274146. 27

[52] R. Subedi et al. **Probing Cold Dense Nuclear Matter**. *Science*, **320**:1476–1478, 2008. doi:10.1126/science.1156675. 28

[53] R. Acciarri et al. **Detection of Back-to-Back Proton Pairs in Charged-Current Neutrino Interactions with the ArgoNeuT Detector in the NuMI Low Energy Beam Line**. *Phys. Rev.*, **D90**(1):012008, 2014. doi:10.1103/PhysRevD.90.012008. 28

[54] L. B. Weinstein, O. Hen, and E. Piasetzky. **"Hammer" events, neutrino energies, and nucleon-nucleon correlations**. *Phys. Rev.*, **C94**(4):045501, 2016. doi:10.1103/PhysRevC.94.045501. 28

[55] C. Ciofi degli Atti and S. Simula. **Realistic model of the nucleon spectral function in few- and many-nucleon systems**. *Phys. Rev. C*, **53**:1689–1710, Apr 1996. doi:10.1103/PhysRevC.53.1689. 28

[56] O. Benhar, A. Fabrocini, S. Fantoni, and I. Sick. **Spectral function of finite nuclei and scattering of GeV electrons**. *Nucl. Phys.*, **A579**:493–517, 1994. doi:10.1016/0375-9474(94)90920-2. 28, 34

[57] O. Benhar. **Nuclear response beyond the Fermi gas model**. In *2nd International Workshop on Neutrino-Nucleus Interactions in the Few GeV Region (NuInt 02) Irvine, California, December 12-15, 2002*, 2003. 28, 29, 33

[58] A. A. Aguilar-Arevalo et al. **First measurement of the muon neutrino charged current quasielastic double differential cross section**. *Phys. Rev. D*, **81**:092005, May 2010. doi:10.1103/PhysRevD.81.092005. 29, 44

[59] M. Martini, M. Ericson, G. Chanfray, and J. Marteau. **A Unified approach for nucleon knock-out, coherent and incoherent pion production in neutrino interactions with nuclei**. *Phys. Rev.*, **C80**:065501, 2009. doi:10.1103/PhysRevC.80.065501. 29, 30, 31

[60] A. Bodek, H. S. Budd, and M. E. Christy. **Neutrino Quasielastic Scattering on Nuclear Targets: Parametrizing Transverse Enhancement (Meson Exchange Currents)**. *Eur. Phys. J.*, **C71**:1726, 2011. doi:10.1140/epjc/s10052-011-1726-y. 32

[61] J. T. Sobczyk. **Transverse Enhancement Model and MiniBooNE Charge Current Quasi-Elastic Neutrino Scattering Data**. *Eur. Phys. J.*, **C72**:1850, 2012. doi:10.1140/epjc/s10052-011-1850-8. 32

[62] J. E. Amaro, M. B. Barbaro, J. A. Caballero, T. W. Donnelly, A. Molinari, and I. Sick. **Using electron scattering superscaling to predict charge-changing neutrino cross sections in nuclei**. *Phys. Rev. C*, **71**:015501, Jan 2005. doi:10.1103/PhysRevC.71.015501. 32

[63] T. Katori. **Meson Exchange Current (MEC) Models in Neutrino Interaction Generators**. *AIP Conf. Proc.*, **1663**(1):030001, 2015. doi:10.1063/1.4919465. 32

[64] M. C. Martinez, P. Lava, N. Jachowicz, J. Ryckebusch, K. Vantournhout, and J. M. Udias. **Relativistic models for quasielastic neutrino scattering**. *Phys. Rev.*, **C73**:024607, 2006. doi:10.1103/PhysRevC.73.024607. 33, 34

[65] J. Ryckebusch, D. Debruyne, P. Lava, S. Janssen, B. Van Overmeire, and T. Van Cauteren. **Relativistic formulation of Glauber theory for A(e, e-prime p) reactions**. *Nucl. Phys.*, **A728**:226–250, 2003. doi:10.1016/j.nuclphysa.2003.08.022. 34

[66] T. Golan, C. Juszczak, and J. T. Sobczyk. **Final State Interactions Effects in Neutrino-Nucleus Interactions**. *Phys. Rev.*, **C86**:015505, 2012. doi:10.1103/PhysRevC.86.015505. 34

[67] T. Leitner, L. Alvarez-Ruso, and U. Mosel. **Charged current neutrino nucleus interactions at intermediate energies**. *Phys. Rev.*, **C73**:065502, 2006. doi:10.1103/PhysRevC.73.065502. 34

[68] S. Boyd, S. Dytman, E. Hernández, J. Sobczyk, and R. Tacik. **Comparison of Models of Neutrino-Nucleus Interactions**. *AIP Conference Proceedings*, **1189**(1):60–73, 2009. doi:10.1063/1.3274191. 35

[69] S. Boyd, S. Dytman, E. Hernandez, J. Sobczyk, and R. Tacik. **Comparison of models of neutrino-nucleus interactions**. *AIP Conf. Proc.*, **1189**(1):60–73, 2009. doi:10.1063/1.3274191. 35

[70] J. E. Sobczyk. **Intercomparison of lepton-nucleus scattering models in the quasielastic region**. *Phys. Rev.*, **C96**(4):045501, 2017. doi:10.1103/PhysRevC.96.045501. 35, 39

[71] K. Abe et al. **Observation of Electron Neutrino Appearance in a Muon Neutrino Beam**. *Phys. Rev. Lett.*, **112**:061802, 2014. doi:10.1103/PhysRevLett.112.061802. 36

[72] T. Nosek. **Systematic Uncertainties of the NOvA Neutrino Oscillation Analysis**. *PoS*, **ICHEP2020**:210, 2021. doi:10.22323/1.390.0210. 36

[73] J. Wolcott. **Impact of cross section uncertainties on NOvA oscillation analyses**. *PoS*, **NuFACT2018**:098, 2019. doi:10.22323/1.341.0098. 36

[74] K. Abe et al. **Hyper-Kamiokande Design Report**. 5 2018. 37

[75] R. Acciarri et al. **Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 2: The Physics Program for DUNE at LBNF**. 12 2015. 37

[76] T. Katori and M. Martini. **Neutrino-Nucleus Cross Sections for Oscillation Experiments**. 2016. 39

[77] O. Lalakulich, U. Mosel, and K. Gallmeister. **Energy reconstruction in quasielastic scattering in the MiniBooNE and T2K experiments**. *Phys. Rev.*, **C86**:054606, 2012. doi:10.1103/PhysRevC.86.054606. 39, 40

[78] A. Cheşneanu and A. Jipa. **Neutrino Energy Reconstruction In Neutrino-Nucleus Interactions**. *Romanian Journal of Physics*, **57**(7-8):1079–1087, 2012. 39

[79] T. Leitner and U. Mosel. **Neutrino-nucleus scattering re-examined: Quasielastic scattering and pion production entanglement and implications for neutrino energy reconstruction**. *Phys. Rev.*, **C81**:064614, 2010. doi:10.1103/PhysRevC.81.064614. 39

[80] M. Wascko. **Alternative Cross Section Analysis Methods**. Talk at T2K Pre-Collaboration Cross-Section Workshop, September 2016. 40

[81] R. D. Cousins, S. J. May, and Y. Sun. **Should unfolded histograms be used to test hypotheses?** 7 2016. 40

[82] A. R. Back et al. **Accelerator Neutrino Neutron Interaction Experiment (ANNIE): Preliminary Results and Physics Phase Proposal**. 2017. 41, 50

[83] R. Ransome. **Pion Absorption in Nuclei**. Talk at Institute for Nuclear Theory (INT) Workshop INT-13-54W, Neutrino-Nucleus Interactions for Current and Next Generation Neutrino Oscillation Experiments, December 2013. 41

[84] D. H. Perkins. **Proton Decay Experiments**. *Annual Review of Nuclear and Particle Science*, **34**(1):1–50, 1984. doi:10.1146/annurev.ns.34.120184.000245. 42

[85] H. Nishino et al. **Search for Proton Decay via p —> e+ pi0 and p —> mu+ pi0 in a Large Water Cherenkov Detector**. *Phys. Rev. Lett.*, **102**:141801, 2009. doi:10.1103/PhysRevLett.102.141801. 42

[86] S.-K. C. Abe, K. et. al. **Search for proton decay via $p \to \nu K^+$ using** 260 kiloton·year **data of Super-Kamiokande**. *Phys. Rev. D*, **90**:072005, Oct 2014. doi:10.1103/PhysRevD.90.072005. 42

[87] I. Anghel et al. **Letter of Intent: The Accelerator Neutrino Neutron Interaction Experiment (ANNIE)**. 2015. 42, 73

[88] C. Xu. **Current status of SK-Gd project and EGADS**. *J. Phys. Conf. Ser.*, **718**(6):062070, 2016. doi:10.1088/1742-6596/718/6/062070. 42, 70

[89] L. Alvarez Ruso et al. **Research and Development for Near Detector Systems Towards Long Term Evolution of Ultra-precise Long-baseline Neutrino Experiments**. 2019. 43

[90] H.-F. Tsoi. **Effects of Missing Neutrons on Neutrino Energy Reconstruction**. 2017. Technical report. 43

[91] S. J. Barish, J. Campbell, G. Charlton, Y. Cho, M. Derrick, R. Engelmann, L. G. Hyman, D. Jankowski, A. Mann, B. Musgrave, P. Schreiner, P. F. Schultz, R. Singer, M. Szczekowski, T. Wangler, H. Yuta, V. E. Barnes, D. D. Carmony, A. F. Garfinkel, and G. M. Radecky. **Study of neutrino interactions in hydrogen and deuterium: Description of the experiment and study of the reaction $\nu + d \to \mu^- + p + p_s$**. *Phys. Rev. D*, **16**:3103–3121, Dec 1977. doi:10.1103/PhysRevD.16.3103. 44

[92] N. J. Baker, A. M. Cnops, P. L. Connolly, S. A. Kahn, H. G. Kirk, M. J. Murtagh, R. B. Palmer, N. P. Samios, and M. Tanaka. **Quasielastic neutrino scattering: A measurement of the weak nucleon axial-vector form factor**. *Phys. Rev. D*, **23**:2499–2505, Jun 1981. doi:10.1103/PhysRevD.23.2499. 44

[93] C. Wilkinson, P. Rodrigues, S. Cartwright, L. Thompson, and K. McFarland. **Reanalysis of bubble chamber measurements of muon-neutrino induced single pion production**. *Phys. Rev.*, **D90**(11):112017, 2014. doi:10.1103/PhysRevD.90.112017. 44

[94] Patrignani, C. *et al*(PDG). **Review of Particle Physics**. *Chinese Physics C*, **40**(10):100001, 2016. 44, 46

[95] A. A. Aguilar-Arevalo and others (MiniBooNE Collaboration). **MiniBooNE Publications**. `https://www-boone.fnal.gov/publications/`. Accessed: 2020-03-21. 44

[96] A. A. Aguilar-Arevalo and others (MiniBooNE Collaboration). **MiniBooNE Data Releases**. `https://www-boone.fnal.gov/for_physicists/data_release/`. Accessed: 2020-03-21. 44

[97] K. Abe et al. **Measurement of double-differential muon neutrino charged-current interactions on C_8H_8 without pions in the final state using the T2K off-axis beam**. *Phys. Rev. D*, **93**(11):112012, 2016. `doi:10.1103/PhysRevD.93.112012`. 44

[98] K. Abe et al. **Measurement of inclusive double-differential ν_μ charged-current cross section with improved acceptance in the T2K off-axis near detector**. *Phys. Rev. D*, **98**:012004, 2018. `doi:10.1103/PhysRevD.98.012004`. 44

[99] V. A. Radescu. **Differential cross section results from NuTeV**. In *Deep inelastic scattering. Proceedings, 12th International Workshop, DIS 2004, Strbske Pleso, Slovakia, April 14-18, 2004. Vol. 1 + 2*, pages 386–391, 2004. 44

[100] Anderson, C. *et al*(ArgoNeuT Collaboration). **First Measurements of Inclusive Muon Neutrino Charged Current Differential Cross Sections on Argon**. *Phys. Rev. Lett.*, **108**:161802, Apr 2012. `doi:10.1103/PhysRevLett.108.161802`. 44

[101] Acciarri, R. *et al*(ArgoNeuT Collaboration). **Measurements of inclusive muon neutrino and antineutrino charged current differential cross sections on argon in the NuMI antineutrino beam**. *Phys. Rev. D*, **89**:112003, Jun 2014. `doi:10.1103/PhysRevD.89.112003`. 44

[102] Tice, B. G. *et al*(MINERvA Collaboration). **Measurement of Ratios of ν_μ Charged-Current Cross Sections on C, Fe, and Pb to CH at Neutrino Energies 2˘20 GeV**. *Phys. Rev. Lett.*, **112**:231801, Jun 2014. `doi:10.1103/PhysRevLett.112.231801`. 44

[103] A. V. Butkevich. **Analysis of quasi-elastic neutrino charged-current scattering off O-16 and neutrino energy reconstruction**. *Phys. Rev.*, **C78**:015501, 2008. `doi:10.1103/PhysRevC.78.015501`. 44

[104] T. Leitner, O. Buss, L. Alvarez-Ruso, and U. Mosel. **Electron- and neutrino-nucleus scattering from the quasielastic to the resonance region**. *Phys. Rev.*, **C79**:034601, 2009. `doi:10.1103/PhysRevC.79.034601`. 44

[105] A. V. Butkevich and S. P. Mikheyev. **Test of Fermi gas model and plane-wave impulse approximation against electron-nucleus scattering data**. *Phys. Rev.*, **C72**:025501, 2005. `doi:10.1103/PhysRevC.72.025501`. 44

[106] A. M. Ankowski and O. Benhar. **Neutron knockout in neutral-current neutrino-oxygen interactions**. *Phys. Rev.*, **D88**:093004, 2013. `doi:10.1103/PhysRevD.88.093004`. 45

[107] R. Gran et al. **Measurement of the quasi-elastic axial vector mass in neutrino-oxygen interactions**. *Phys. Rev.*, **D74**:052002, 2006. `doi:10.1103/PhysRevD.74.052002`. 45, 156

[108] Abe, K. *et al*(T2K Collaboration). **Measurement of the neutrino-oxygen neutral-current interaction cross section by observing nuclear deexcitation γ rays**. *Phys. Rev. D*, **90**:072012, Oct 2014. `doi:10.1103/PhysRevD.90.072012`. 45, 156

[109] K. Abe and O. T. Collaboration). **Measurement of neutrino and antineutrino neutral-current quasielasticlike interactions on oxygen by detecting nuclear deexcitation γ rays**. *Phys. Rev. D*, **100**:112009, Dec 2019. `doi:10.1103/PhysRevD.100.112009`. 45, 156

[110] K. Abe and O. T. Collaboration). **First measurement of the ν_μ charged-current cross section on a water target without pions in the final state**. *Phys. Rev. D*, **97**:012001, Jan 2018. `doi:10.1103/PhysRevD.97.012001`. 45, 156

[111] O. Palamara, K. Partyka, and F. Cavanna. **Exclusive Topologies reconstruction in LAr-TPC experiments: a Novel Approach for precise Neutrino-Nucleus Cross-Sections Measurements**. In *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*, 2013. 47

[112] Acciarri, R. *et al*(ArgoNeuT Collaboration). **Detection of back-to-back proton pairs in charged-current neutrino interactions with the ArgoNeuT detector in the NuMI low energy beam line**. *Phys. Rev. D*, **90**:012008, Jul 2014. `doi:10.1103/PhysRevD.90.012008`. 47

[113] R. Acciarri et al. **Demonstration of MeV-Scale Physics in Liquid Argon Time Projection Chambers Using ArgoNeuT**. *Phys. Rev.*, **D99**(1):012002, 2019. `doi:10.1103/PhysRevD.99.012002`. 47

[114] H. Zhang. **Neutron tagging and its application to physics in Super-Kamiokande IV**. In *32nd International Cosmic Ray Conference*, 2011. `doi:10.7529/ICRC2011/V04/0353`. 47

[115] A. Back et al. **Measurement of Beam-Correlated Background Neutrons from the Fermilab Booster Neutrino Beam in ANNIE Phase-I**. *JINST*, **15**(03):P03011, 2020. `doi:10.1088/1748-0221/15/03/P03011`. 50, 121

[116] A. Meyer. **The CDF experiment at the Tevatron: The First two years of Run II**. *Mod. Phys. Lett.*, **A18**:1643–1660, 2003. `doi:10.1142/S0217732303011599`. 52

[117] P. Cherenkov. **"Visible glow of pure liquids under gamma irradiation"**. *Usp. Fiz. Nauk*, **93**(10):385–388, 1967. `doi:10.3367/UFNr.0093.196710n.0385`. 54

[118] I. Frank and I. Tamm. **Coherent visible radiation of fast electrons passing through matter**. *Compt. Rend. Acad. Sci. URSS*, **14**(3):109–114, 1937. `doi:10.3367/UFNr.0093.196710o.0388`. 54

[119] W.-M. Yao and others. Particle Data Group. **Review of Particle Physics**. *Journal of Physics G*, **33**:1+, 2006. 54, 113

[120] F. Montanet. **Cherenkov and Imaging Detectors for HEP and AP**. Talk at The European School in Instrumentation for Particle and Astroparticle Physics (ESIPAP), January 2014. 55

[121] R. Becker-Szendy and others (IMB Collaboration). **IMB-3: a large water Cherenkov detector for nucleon decay and neutrino interactions**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **324**(1):363 – 382, 1993. `doi:https://doi.org/10.1016/0168-9002(93)90998-W`. 55

[122] J. Boger et al. **The Sudbury neutrino observatory**. *Nucl. Instrum. Meth.*, **A449**:172–207, 2000. `doi:10.1016/S0168-9002(99)01469-2`. 55

[123] K. Hirata et al. **Observation in the Kamiokande-II Detector of the Neutrino Burst from Supernova SN 1987a**. *Phys. Rev. D*, **38**:448–458, 1988. `doi:10.1103/PhysRevD.38.448`. 55

[124] Y. Fukuda et al. **The Super-Kamiokande detector**. *Nucl. Instrum. Meth. A*, **501**:418–462, 2003. doi:10.1016/S0168-9002(03)00425-X. 55

[125] N. Kurahashi. **The Birth of Neutrino Astronomy**. In *Proceedings, 10th International Symposium on Cosmology and Particle Astrophysics (CosPA 2013): Honolulu, Hawaii, USA, November 12-15, 2013*, 2014. 55

[126] J. Aublin, G. Illuminati, and S. Navas. **Searches for point-like sources of cosmic neutrinos with 11 years of ANTARES data**. *PoS*, **ICRC2019**:920, 2020. doi:10.22323/1.358.0920. 55

[127] K. Abe et al. **Hyper-Kamiokande Design Report**. 2018. 55

[128] S. Adrian-Martinez et al. **Letter of intent for KM3NeT 2.0**. *J. Phys.*, **G43**(8):084001, 2016. doi:10.1088/0954-3899/43/8/084001. 55

[129] **Kentain Products Ltd**. https://www.kentain.com. Accessed: 2020-03-21. 56

[130] Y. Nishimura. **New 50cm Photo-Detectors for Hyper-Kamiokande**. Talk at the 2016 International Conference on High-Energy Physics (ICHEP), August 2016. 57

[131] A. Rohatgi. **Webplotdigitizer: Version 4.3**, 2020. 59

[132] D. S. Akerib et al. **The Large Underground Xenon (LUX) Experiment**. *Nucl. Instrum. Meth.*, **A704**:111–126, 2013. doi:10.1016/j.nima.2012.11.135. 57

[133] N. Barros, T. Kaptanoglu, B. Kimelman, J. R. Klein, E. Moore, J. Nguyen, K. Stavreva, and R. Svoboda. **Characterization of the ETEL D784UKFLB 11 in. photomultiplier tube**. *Nucl. Instrum. Meth.*, **A852**:15–19, 2017. doi:10.1016/j.nima.2017.01.067. 57

[134] S. Dazeley, M. Askins, M. Bergevin, A. Bernstein, N. S. Bowden, T. M. Shokair, P. Jaffke, S. D. Rountree, and M. Sweany. **A search for cosmogenic production of β-neutron emitting radionuclides in water**. *Nucl. Instrum. Meth.*, **A821**:151–159, 2016. doi:10.1016/j.nima.2016.03.014. 58

[135] M. Askins et al. **The Physics and Nuclear Nonproliferation Goals of WATCHMAN: A WAter CHerenkov Monitor for ANtineutrinos**. 2015. 58

[136] Y. Nakajima. **Status of FNAL SciBooNE experiment**. *J. Phys. Conf. Ser.*, **120**:052043, 2008. doi:10.1088/1742-6596/120/5/052043. 62

[137] R. Ben-David. **Status of the KTeV experiment at Fermilab**. *Nucl. Phys. Proc. Suppl.*, **66**:473–477, 1998. doi:10.1016/S0920-5632(98)00088-7. 62

[138] Y. Sugiyama et al. **The Data Acquisition System for the KOTO Experiment**. *IEEE Trans. Nucl. Sci.*, **62**(3):1115–1121, 2015. doi:10.1109/RTC.2014.7097532,10.1109/TNS.2015.2417312. 66

[139] B. Richards. **ToolDAQ github page**. https://github.com/ToolDAQ/ToolDAQFramework. Accessed: 2020-03-21. doi:https://doi.org/10.5281/zenodo.3229251. 68

[140] Richards, Benjamin. **The ToolDAQ DAQ Software Framework & Its Use In The Hyper-K & ANNIE Detectors**. *EPJ Web Conf.*, **214**:01022, 2019. doi:10.1051/epjconf/201921401022. 68

[141] **slack**. https://slack.com/. Accessed: 2020-03-21. 69

[142] H. Zhang. **Neutron tagging and its physics application in Super-Kamiokande-IV**. *Proceedings of the 32nd International Cosmic Ray Conference, ICRC 2011*, **4**:71–, 01 2011. doi:10.7529/ICRC2011/V04/0353. 70

[143] J. F. Beacom and M. R. Vagins. **GADZOOKS! Anti-neutrino spectroscopy with large water Cherenkov detectors**. *Phys. Rev. Lett.*, **93**:171101, 2004. doi:10.1103/PhysRevLett.93.171101. 70

[144] B. Baramsai. *Neutron Capture Reactions on Gadolinium Isotopes.* PhD thesis, February 2010. 70

[145] K. Hagiwara et al. **Gamma Ray Spectrum from Thermal Neutron Capture on Gadolinium-157**. *PTEP*, **2019**(2):023D01, 2019. doi:10.1093/ptep/ptz002. 70, 71, 72

[146] T. Tanaka, K. Hagiwara, E. Gazzola, A. Ali, I. Ou, T. Sudo, P. Kumar Das, M. Singh Reen, R. Dhir, Y. Koshio, M. Sakuda, A. Kimura, S. Nakamura, N. Iwamoto, H. Harada, G. Collazuol, S. Lorenz, M. Wurm, W. Focillon, M. Gonin, and T. Yano. **Gamma-ray spectra from thermal neutron capture on gadolinium-155 and natural gadolinium**. *Progress of Theoretical and Experimental Physics*, **2020**(4), 04 2020. 043D02. doi:10.1093/ptep/ptaa015. 70

[147] D. Castelluccio, A. Mengoni, and A. Guglielmelli. **Implementation of a cross section evaluation methodology for safety margins analysis: Post processing and data analysis for the capture cross section measurements of odd Gadolinium isotopes**. Technical report, 11 2018. doi:10.13140/RG.2.2.14596.17283. 70

[148] D. Brown and Others. **ENDF/B-VIII.0: The 8th Major Release of the Nuclear Reaction Data Library with CIELO-project Cross Sections, New Standards and Thermal Scattering Data**. *Nuclear Data Sheets*, **148**:1 – 142, 2018. Special Issue on Nuclear Reaction Data. doi:https://doi.org/10.1016/j.nds.2018.02.001. 71

[149] T. Mori. **Development of A Gadolinium-doped Water Cherenkov Detector for The Observation of Supernova Relic Neutrinos**. Talk at Institute for Nuclear Theory (INT) Workshop INT-13-54W, Neutrino-Nucleus Interactions for Current and Next Generation Neutrino Oscillation Experiments, 2015. 72

[150] H. Watanabe et al. **First Study of Neutron Tagging with a Water Cherenkov Detector**. *Astropart. Phys.*, **31**:320–328, 2009. doi:10.1016/j.astropartphys.2009.03.002. 71

[151] A. Renshaw. **Research and Development for a Gadolinium Doped Water Cherenkov Detector**. *Phys. Procedia*, **37**:1249–1256, 2012. doi:10.1016/j.phpro.2012.02.467. 73, 78

[152] R. Dharmapalan et al. **A new investigation of electron neutrino appearance oscillations with improved sensitivity in the MiniBooNE+ experiment**. In *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*, 2013. 73

[153] C. Blanco. **ANNIE and the Characterization of CCQE Backgrounds in Genie MC**. Talk at the Final State Nucleons Meeting, Fermilab, August 2017. 73

[154] P. Fernandez. **Gd-doping and the impact on SuperK and T2K**. Talk at the International Workshop on Next Generation Nucleon Decay and Neutrino Detectors (NNN16), August 2016. 73

[155] **EMI Electro-Max Inc.** https://www.electromax-inc.com/. Accessed: 2020-03-21. 74

[156] T. G. Gooch. **Corrosion behavior of welded stainless steel**. *Welding Journal*, **75**(5), 5 1996. 74

[157] J. He. **UC Davis Material Compatibility Studies**. http://svoboda.ucdavis.edu/experiments/watchman/materials-testing/, Accessed: 2020-10-23. 76

[158] L. Marti et al. **Evaluation of Gadolinium's Action on Water Cherenkov Detector Systems with EGADS**. *Nucl. Instrum. Meth.*, **A959**:163549, 2020. doi:10.1016/j.nima.2020.163549. 78

[159] V. Fischer, J. He, M. Irving, and R. Svoboda. **Development of an ion exchange resin for gadolinium-loaded water**. *JINST*, **15**(07):P07004, 2020. doi:10.1088/1748-0221/15/07/P07004. 78

[160] P. Vogel and J. F. Beacom. **Angular distribution of neutron inverse beta decay, anti-neutrino(e) + p —¿ e+ + n**. *Phys. Rev.*, **D60**:053003, 1999. doi:10.1103/PhysRevD.60.053003. 82

[161] B. W. Adams et al. **A Brief Technical History of the Large-Area Picosecond Photodetector (LAPPD) Collaboration**. *Submitted to: JINST*, 2016. 82

[162] J. Xie. **Characteristics of MCP-PMTs/LAPPDs in Magnetic Field**. Talk at the 39th International Conference On High Energy Physics (ICHEP2018), July 2018. 84

[163] J. Anderson et al. **Considerations about Large Area–Low Cost Fast Imaging Photo-detectors**. In *Proceedings, 2011 Acta Physica Polonica B*, 2011. doi:10.5506/APhysPolBSupp.4.13. 84

[164] H. Grabas, R. Obaid, E. Oberla, H. Frisch, J.-F. Genat, R. Northrop, F. Tang, D. McGinnis, B. Adams, and M. Wetstein. **RF strip-line anodes for Psec large-area MCP-based photodetectors**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **711**:124 – 131, 2013. doi:https://doi.org/10.1016/j.nima.2013.01.055. 86

[165] M. Popecki. **The Large Area Picosecond Photodetector (LAPPD) 8" MCP-PMT: Recent Results**. Research Techniques Seminar, Fermilab, March 2019. 86, 87

[166] E. Oberla, J.-F. Genat, H. Grabas, H. Frisch, K. Nishimura, and G. Varner. **A 15 GSa/s, 1.5 GHz bandwidth waveform digitizing ASIC**. *Nucl. Instrum. Meth.*, **A735**:452–461, 2014. doi:10.1016/j.nima.2013.09.042. 87, 88

[167] J.-F. Genat, G. Varner, F. Tang, and H. J. Frisch. **Signal Processing for Pico-second Resolution Timing Measurements**. *Nucl. Instrum. Meth.*, **A607**:387–393, 2009. doi:10.1016/j.nima.2009.05.193. 87

[168] G. R. Jocher, M. J. Wetstein, B. Adams, K. Nishimura, and S. M. Usman. **Multiple-photon disambiguation on stripline-anode Micro-Channel Plates**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **822**:25 – 33, 2016. doi:https://doi.org/10.1016/j.nima.2016.03.079. 87, 153

[169] F. Kaether and C. Langbrandtner. **Transit time and charge correlations of single photoelectron events in R7081 photomultiplier tubes**. *Journal of Instrumentation*, **7**(09):P09002–P09002, sep 2012. doi:10.1088/1748-0221/7/09/p09002. 91

[170] P. A. Amaudruz et al. **In-situ characterization of the Hamamatsu R5912-HQE photomultiplier tubes used in the DEAP-3600 experiment**. *Nucl. Instrum. Meth.*, **A922**:373–384, 2019. doi:10.1016/j.nima.2018.12.058. 91, 92

[171] R. Saldanha, L. Grandi, Y. Guardincerri, and T. Wester. **Model Independent Approach to the Single Photoelectron Calibration of Photomultiplier Tubes**. *Nucl. Instrum. Meth.*, **A863**:35–46, 2017. doi:10.1016/j.nima.2017.02.086. 91

[172] K. Ma, W. Kang, J. Ahn, S. Choi, Y. Choi, M. Hwang, J. Jang, E. Jeon, K. Joo, H. Kim, and et al. **Time and amplitude of afterpulse measured with a large size photomultiplier tube**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **629**(1):93–100, Feb 2011. doi:10.1016/j.nima.2010.11.095. 93

[173] J. Haser, F. Kaether, C. Langbrandtner, M. Lindner, S. Lucht, S. Roth, M. Schumann, A. Stahl, A. Stüken, and C. Wiebusch. **Afterpulse Measurements of R7081 Photomultipliers for the Double Chooz Experiment**. *JINST*, **8**:P04029, 2013. doi:10.1088/1748-0221/8/04/P04029. 93

[174] **HVSys Company**. http://hvsys.ru/. Accessed: 2020-03-21. 93

[175] **Exact Metrology**. https://www.exactmetrology.com/. Accessed: 2020-03-21. 94

[176] T. J. Pershing. *The Accelerator Neutrino-Neutron Interaction Experiment*. PhD thesis, UC, Davis, 2020. doi:10.2172/1638651. 97, 125

[177] E. Tiras. **Detector R&D for ANNIE and Future Neutrino Experiments**. In *Meeting of the Division of Particles and Fields of the American Physical Society*, 10 2019. 99

[178] J. Crawford et al. **Fermilab Operations Booster Rookie Book**. https://operations.fnal.gov/rookie_books/Booster_V4.1.pdf. Accessed: 2020-03-21. 105

[179] M. Gattuso. **MiniBooNE Rookie Book**. https://operations.fnal.gov/rookie_books/MiniBooNE_v1.pdf. Accessed: 2020-03-21. 105

[180] **Fermilab Accelerator Complex web site**. https://www.fnal.gov/pub/science/particle-accelerators/accelerator-complex.html. Accessed: 2020-03-21. 105

[181] A. Aguilar-Arevalo et al. **The Neutrino Flux prediction at MiniBooNE**. *Phys. Rev. D*, **79**:072002, 2009. doi:10.1103/PhysRevD.79.072002. 106, 109, 110, 113, 114, 115

[182] I. Stancu and B. C. others. **Technical Design Report for the MiniBooNE Neutrino Beam**. https://www-boone.fnal.gov/publicpages/target_tdr.ps.gz. Accessed: 2020-03-21. 106

[183] **Fermilab Target Systems BNB web site**. https://targets.fnal.gov/BNB_neutrino_beam.html. Accessed: 2020-03-21. 107

[184] **Fermilab Target System Facilities web site**. https://targets.fnal.gov/TSD_facilities.html. Accessed: 2020-10-23. 107

[185] D. W. Schmitz. *A Measurement of Hadron Production Cross Sections for the Simulation of Accelerator Neutrino Beams and a Search for ν_μ to ν_e Sscillations in the δm^2 about equals $1-eV^2$ Region*. PhD thesis, Columbia U., 2008. doi:10.2172/935240. 109

[186] R. J. Glauber. **High Energy Collision Theory**. **1**:315, 1959. 109

[187] V. Franco. **High-energy neutron-nuclei total cross-sections**. *Phys. Rev. C*, **6**:748–757, 1972. doi:10.1103/PhysRevC.6.748. 110

[188] R. T. **Introduction to complex orbital momenta**. *Nuovo Cim*, **14**:951–976, 1959. doi:10.1007/BF02728177. 110

[189] M. Wilking. *Measurement of neutrino induced, charged current, charged pion production.* PhD thesis, 5 2009. doi:10.2172/1369266. 111, 112, 114

[190] J. Sanford and C. Wang. **Empirical Formulas for Particle Production in P-Be Collisions between 10 and 35 GeV/c**. Brookhaven National Laboratory Reports. BNL 11299 and BNL 11479, March 1967. 110

[191] S. Brice. **An Improved Way to Propagate Pi+ Production Uncertainties**. MiniBooNE Technical Note 257, 2008. 110

[192] G. C. L. Cheng. *Precision Search for Muon Antineutrino Disappearance Oscillations Using a Dual Baseline Technique.* PhD thesis, Columbia U., 2013. doi:10.2172/1346824. 110

[193] K. Hiraide and Others. **Search for charged current coherent pion production on carbon in a few-GeV neutrino beam**. *Phys. Rev. D*, **78**:112004, Dec 2008. doi:10.1103/PhysRevD.78.112004. 115

[194] C. Andreopoulos et al. **The GENIE Neutrino Monte Carlo Generator**. *Nucl. Instrum. Meth. A*, **614**:87–104, 2010. doi:10.1016/j.nima.2009.12.009. 116

[195] A. Bodek and J. L. Ritchie. **Fermi-motion effects in deep-inelastic lepton scattering from nuclear targets**. *Phys. Rev. D*, **23**:1070–1091, Mar 1981. doi:10.1103/PhysRevD.23.1070. 116

[196] A. Bodek and J. L. Ritchie. **Further studies of Fermi-motion effects in lepton scattering from nuclear targets**. *Phys. Rev. D*, **24**:1400–1402, Sep 1981. doi:10.1103/PhysRevD.24.1400. 116

[197] A. Bodek, S. Avvakumov, R. Bradford, and H. S. Budd. **Vector and Axial Nucleon Form Factors:A Duality Constrained Parameterization**. *Eur. Phys. J. C*, **53**:349–354, 2008. doi:10.1140/epjc/s10052-007-0491-4. 116

[198] T. Yang, C. Andreopoulos, H. Gallagher, K. Hoffmann, and P. Kehayias. **A Hadronization Model for Few-GeV Neutrino Interactions**. *Eur. Phys. J. C*, **63**:1–10, 2009. doi:10.1140/epjc/s10052-009-1094-z. 117

[199] T. Sjöstrand, S. Mrenna, and P. Skands. **A brief introduction to PYTHIA 8.1**. *Computer Physics Communications*, **178**(11):852 – 867, 2008. doi:https://doi.org/10.1016/j.cpc.2008.01.036. 117

[200] A. M. Ankowski and O. Benhar. **Neutron knockout in neutral-current neutrino-oxygen interactions**. *Phys. Rev. D*, **88**:093004, 2013. doi:10.1103/PhysRevD.88.093004. 117

[201] A. Rindi and R. Thomas. **Skyshine : A Paper Tiger?** *Part. Accel.*, **7**:23–39, 1975. 118

[202] S. Lindenbaum. **Skyshine**. In *Proceedings of the Conference on Shielding of High-Energy Accelerators*, pages 101–111, 1957. doi:10.1016/0370-1573(72)90010-5. 118

[203] A. Aguilar-Arevalo et al. **Bringing the SciBar detector to the booster neutrino beam**. 1 2006. 119

[204] S. J. Gardiner. *Nuclear Effects in Neutrino Detection.* PhD thesis, 2018. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2019-10-18. 121

[205] S. Agostinelli and Others. **Geant4—a simulation toolkit**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **506**(3):250 – 303, 2003. doi:https://doi.org/10.1016/S0168-9002(03)01368-8. 123

[206] **WCSim github page**. https://github.com/WCSim/WCSim. Accessed: 2020-03-21. 123

[207] J. Allison and Others. **Recent developments in Geant4**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **835**:186 – 225, 2016. doi:https://doi.org/10.1016/j.nima.2016.06.125. 123

[208] M. Chadwick and Others. **ENDF/B-VII.1 Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data**. *Nuclear Data Sheets*, **112**(12):2887 – 2996, 2011. Special Issue on ENDF/B-VII.1 Library. doi:https://doi.org/10.1016/j.nds.2011.11.002. 123

[209] R. C. E. Mendoza, D. Cano-Ott. **Update of the Evaluated Neutron Cross Section Libraries for the Geant4 Code**. June 2018. Accessed: 2020-03-21. 123

[210] C. G. E. Mendoza, D. Cano-Ott and R. Capote. **New Evaluated Neutron Cross Section Libraries for the Geant4 Code**. April 2012. Accessed: 2020-03-21. 123

[211] A. Levin and C. Moisan. **A more physical approach to model the surface treatment of scintillation counters and its implementation into DETECT**. In *1996 IEEE Nuclear Science Symposium. Conference Record*, **2**, pages 702–706 vol.2, 1996. 123

[212] F. James and M. Roos. **Minuit - a system for function minimization and analysis of the parameter errors and correlations**. *Computer Physics Communications*, **10**(6):343 – 367, 1975. doi:https://doi.org/10.1016/0010-4655(75)90039-9. 131

[213] H. Maesaka. *Evidence for Muon Neutrino Oscillation in an Accelerator Based Experiment.* PhD thesis, Kyoto U., 2005. 137

[214] T. Kikawa. *Measurement of Neutrino Interactions and Three Flavor Neutrino Oscillations in the T2K Experiment.* PhD thesis, Kyoto U. (main), 9 2014. doi:10.1007/978-981-287-715-4. 137

[215] M. Abadi et al. **TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems**, 2015. Software available from tensorflow.org. 141

[216] J. H. Friedman. **Greedy Function Approximation: A Gradient Boosting Machine**. *The Annals of Statistics*, **29**(5):1189–1232, 2001. 141

[217] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. **Scikit-learn: Machine Learning in Python** . *Journal of Machine Learning Research*, **12**:2825–2830, 2011. 141

[218] M. Shiozawa. **Reconstruction algorithms in the Super-Kamiokande large water Cherenkov detector**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **433**(1):240 – 246, 1999. doi:https://doi.org/10.1016/S0168-9002(99)00359-9. 143

[219] P. Gloor. *New Approach for a Circle Hough Transform to Detect Cherenkov Rings in the LHCb RICH Detector.* PhD thesis, 2016. Accessed: 2020-03-21. 143

[220] B. S. Wonsak, C. I. Hagner, D. A. Hellgartner, K. Loo, S. Lorenz, D. J. Meyhöfer, L. Oberauer, H. Rebber, W. H. Trzaska, and M. Wurm. **Topological track reconstruction in unsegmented, large-volume liquid scintillator detectors**. *JINST*, **13**(07):P07005, 2018. doi:10.1088/1748-0221/13/07/P07005. 144

[221] M. Wetstein. **LAPPDs in Water and Scintillator Detectors**. Talk at the Frontiers in liquid Scintillator Technology (FRoST) 2017 Workshop, March 2017. 145

[222] M. Yeh, S. Hans, W. Beriguete, R. Rosero, L. Hu, R. Hahn, M. Diwan, D. Jaffe, S. Kettell, and L. Littenberg. **A new water-based liquid scintillator and potential applications**. *Nucl. Instrum. Meth. A*, **660**:51–56, 2011. `doi:10.1016/j.nima.2011.08.040`. 154

[223] L. J. Bignell et al. **Characterization and Modeling of a Water-based Liquid Scintillator**. *JINST*, **10**(12):P12009, 2015. `doi:10.1088/1748-0221/10/12/P12009`. 154

[224] J. Alonso et al. **Advanced Scintillator Detector Concept (ASDC): A Concept Paper on the Physics Potential of Water-Based Liquid Scintillator**. 9 2014. 154

[225] J. Alonso et al. **Advanced Scintillator Detector Concept (ASDC): A Concept Paper on the Physics Potential of Water-Based Liquid Scintillator**. 9 2014. 154

[226] C. Aberle, A. Elagin, H. J. Frisch, M. Wetstein, and L. Winslow. **Measuring directionality in double-beta decay and neutrino interactions with kiloton-scale scintillation detectors**. *Journal of Instrumentation*, **9**(06):P06012–P06012, jun 2014. `doi:10.1088/1748-0221/9/06/p06012`. 154

[227] M. Askins et al. **Theia: An advanced optical neutrino detector**. 11 2019. 154

[228] T. Stern, I. Kaner, N. Laser Zer, H. Shoval, D. Dror, Z. Manevitch, L. Chai, Y. Brill-Karniely, and O. Benny. **Rigidity of polymer micelles affects interactions with tumor cells**. *Journal of Controlled Release*, **257**:40 – 50, 2017. New Horizons in Drug and Gene Delivery - ICRS 2016. `doi:https://doi.org/10.1016/j.jconrel.2016.12.013`. 155

[229] E. Angelico, J. Eisch, A. Elagin, H. Frisch, S. Nagaitsev, and M. Wetstein. **Energy and Flavor Discrimination Using Precision Time Structure in On-Axis Neutrino Beams**. *Phys. Rev. D*, **100**(3):032008, 2019. `doi:10.1103/PhysRevD.100.032008`. 155

[230] **Workshop on Precision Time Structure in On-Axis Neutrino Beams**. `https://indico.fnal.gov/event/21409/`. Accessed: 2020-03-21. 155